SEVENTH FRAMEWORK PROGRAMME
Research Infrastructures

INFRA-2010-2.3.1 – First Implementation Phase of the European High Performance Computing (HPC) service PRACE

PRACE-1IP

PRACE First Implementation Project

Grant Agreement Number: RI-261557

D7.4.3
Tier-0 Applications and Systems Usage

Final

Version: 1.0
Author(s): Xu Guo, Mark Bull (EPCC)
Date: 24.05.2012

Project and Deliverable Information Sheet

PRACE Project	**Project Ref. №: RI-261557**	
	Project Title: PRACE First Implementation Project	
	Project Web Site: http://www.prace-project.eu	
	Deliverable ID: D7.4.3	
	Deliverable Nature: DOC_TYPE: Report	
	Deliverable Level: PU	**Contractual Date of Delivery:** 31 / May / 2012
		Actual Date of Delivery: 31 / May /2012
	EC Project Officer: Bernhard Fabianek	

- - The dissemination level are indicated as follows: **PU** – Public, **PP** – Restricted to other participants (including the Commission Services), **RE** – Restricted to a group specified by the consortium (including the Commission Services). **CO** – Confidential, only for members of the consortium (including the Commission Services).

Document Control Sheet

Document	**Title: Tier-0 Applications and Systems usage**	
	ID: D7.4.3	
	Version: 1.0	**Status:**
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2010	
	File(s): D7.4.3-ForApproval.docx	
Authorship	**Written by:**	Xu Guo, EPCC Mark Bull, EPCC
	Contributors:	Stefanie Janetzko, FZJ Patrice Lucas, CEA
	Reviewed by:	Dietmar Erwin, FZJ; Miroslaw Kupczyk, PSNC.
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	20/04/2012	Draft	
0.2	26/04/2012	Draft	Minor changes for figures
0.3	03/05/2012	For internal review	Revised, added Exec Summary and Conclusions
1.0	24/05/2012	For Approval	Revised as requested by internal reviwers.

Document Keywords

Keywords:	PRACE, HPC, Research Infrastructure, Systems, Applications, Users, Surveys

Table of Contents

List of Figures

List of Tables

References and Applicable Documents

[1] PRACE Project: http://www.prace-project.eu
[2] JUGENE website: http://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/JUGENE/Configuration/Configuration_node.html
[3] CURIE website: http://www-hpc.cea.fr/en/complexe/tgcc-curie.htm
[4] PRACE Preparatory Phase Deliverable D6.1 “Identification and categorisation of applications and initial benchmarks suite”.
[5] PRACE First Implementation Phase Deliverable D7.4.1 “Applications and user requirements for Tier-0 systems”.

List of Acronyms and Abbreviations

BSC	Barcelona Supercomputing Center (Spain)
CEA	Commissariat à l’Energie Atomique (represented in PRACE by GENCI, France)
CINECA	Consorzio Interuniversitario, the largest Italian computing centre (Italy).
CSC	Finnish IT Centre for Science (Finland)
DEISA	Distributed European Infrastructure for Supercomputing Applications. EU project by leading national HPC centres.
EPCC	Edinburgh Parallel Computing Centre (represented in PRACE by EPSRC, United Kingdom)
FZJ	Forschungszentrum Jülich (Germany)
GB	Giga (= 2^{30} ~ 10^9) Bytes (= 8 bits), also GByte
GB/s	Giga (= 10^9) Bytes (= 8 bits) per second, also GByte/s
GFlop/s	Giga (= 10^9) Floating point operations (usually in 64-bit, i.e. DP) per second, also GF/s
GHz	Giga (= 10^9) Hertz, frequency =10^9 periods or clock cycles per second
GPGPU	General Purpose GPU
GPU	Graphic Processing Unit
HLRS	High Performance Computing Center Stuttgart (Germany)
HPC	High Performance Computing; Computing at a high performance level at any given time; often used synonym with Supercomputing
HPF	High Performance Fortran
HP-SEE	High-Performance Computing Infrastructure for South East Europe’s Research Communities, Project Acronym
ICHEC	Irish Centre for High-End Computing (Ireland)
IDRIS	Institut du Développement et des Ressources en Informatique Scientifique (represented in PRACE by GENCI, France)
IPB	Institute of Physics, Belgrade (Serbia)
KTH	Kungliga Tekniska Högskolan (represented in PRACE by SNIC, Sweden)
LEF	LINPACK Equivalent Flop/s
LINPACK	Software library for Linear Algebra
LRZ	Leibniz Supercomputing Centre (Garching, Germany)
MB	Mega (= 2^{20} ~ 10^6) Bytes (= 8 bits), also MByte
MB/s	Mega (= 10^6) Bytes (= 8 bits) per second, also MByte/s
MPI	Message Passing Interface
NTNU	Norwegian University of Science and Technology (Trondheim, Norway)
OpenMP	Open Multi-Processing

PFlop/s	Peta (= 10^{15}) Floating-point operations (usually in 64-bit, i.e. DP) per second, also PF/s
PGAS	Partitioned Global Address Space
PLEF	Peta-LEF
PRACE	Partnership for Advanced Computing in Europe; Project Acronym
PRACE-PP	PRACE Preparatory Phase Project
PSNC	Poznan Supercomputing and Networking Centre (Poland)
SEE	South Eastern Europe
SNIC	Swedish National Infrastructure for Computing (Sweden)
TFlop/s	Tera (= 10^{12}) Floating-point operations (usually in 64-bit, i.e. DP) per second, also TF/s
Tier-0	Denotes the apex of a conceptual pyramid of HPC systems. In this context the Supercomputing Research Infrastructure would host the Tier-0 systems; national or topical HPC centres would constitute Tier-1
TLEF	Tera-LEF
UiB	University of Bergen (Norway)
UPC	Unified Parallel C
VRC	Virtual Research Community

Executive Summary

This is the final deliverable for Task 7.4 (Applications Requirements for Tier-0 Systems). It reports the results of surveys carried out of PRACE usage of Tier-0 systems, and of users participating in PRACE Preparatory and Regular Access projects.

The key findings of the surveys are:

- System usage data was collected from JUGENE (BlueGene/P system at FZJ) and CURIE (Bull system at CEA).
- PRACE usage accounts for just over one third of all the CPU hours consumed on JUGENE. This matches the commitment of one third of JUGENE capacity to PRACE.
- Completed Regular Access projects have used all their allocated CPU hours.
- Just over 50% of the cycles used by Regular Access projects on JUGENE were in jobs running on more than 8192 cores.
- 45% of the cycles used by Regular Access projects on CURIE were in jobs running on more than 2048 cores.
- A small number of projects account for a large fraction of the disk usage on both systems.
- Applications exhibit a wide range of disk requirements for running production jobs.
- 62 users from 18 PRACE partner counties responded to the user survey, with users of both systems and all types of project well represented.
- 40 different application codes were named as being used.
- Of these application codes nearly 50% use combined MPI+OpenMP as the parallelisation method. (The remainder use MPI only: no use was reported of other parallelization models such as PGAS.)
- There is a roughly equal split between Fortran and C/C++ as the main language used by applications.
- Higher peak flop rate is seen as the architectural feature that would most benefit applications.
- Over 60% of the applications were reported as having an existing accelerator port, or would be likely to benefit from such.
- In terms of network requirements between PRACE systems and other sites, both capability and capacity are seen as important.
- Only 13% of users know how to use the dedicated PRACE network.
- There is significant interest in having various types of network information available to users.
- Use of, and interest in, Grid technologies is low.
- A significant minority of users is interested in, or makes use of, remote visualization facilities.

1 Introduction

This document contains the results of the Tier-0 system and user surveys that were carried out for the Task 7.4 (Applications Requirements for Tier-0 Systems) of the PRACE First Implementation Phase Project (PRACE-1IP) [1]. The surveys include:

- A survey of the PRACE Tier-0 systems, JUGENE and CURIE
- A survey of the PRACE Tier-0 Access users, including both Preparatory Access and Regular Access

The purposes of these surveys are:

- To understand the current usage status of the PRACE Tier-0 systems.
- To understand the users' experience and further requirements for the PRACE Tier-0 systems

Section 2 of this document summarises the results of the system survey and Section 3 contains the summaries of the user survey responses. Section 0 contains some conclusions.

A number of PRACE-1IP work packages and tasks contributed questions to the user survey, including WP6 (Technical Operation and Evolution of the Distributed Infrastructure) and Task 7.1 (Applications Enabling for Capability Science). The data from the user survey results is expected to be used by the contributing work packages/tasks for more detailed analysis.

2 System Survey

2.1 Overview

The questions of the system survey were asked to understand the real usage status of the PRACE Tier-0 systems. This survey was answered for the Tier-0 BlueGene/P system JUGENE and the Tier-0 Bull x86 cluster system CURIE. The survey was not completed for the Cray XE6 Tier-0 system HERMIT, as HERMIT was not in production status during this survey time.

The survey response for JUGENE included the usage status between 1st June 2010 and 31st December 2011. The response for CURIE included the usage status between 1st March 2011 and 14th March 2012. The answers to the survey questions were collected using a text questionnaire that was completed for both systems by staff at the hosting sites.

Note some PRACE Access projects had not been finished by the end date of survey periods. Timescales of the Regular Access projects on JUGENE are listed below:

- prapro01-prapro19[1]: finished;
- prapro20-prapro29: started on 1st May 2011 and will end on 30th April 2012;
- prapro30-prapro39: started on 1st November 2011 and will end on 31st August 2012.

On CURIE, 59 Preparatory Access projects have finished while 34 Preparatory Access projects were still on-going by the end date of survey period. 18 Regular Access projects are all due on 1st November 2012.

The system survey concentrated on the following aspects:

[1] Note that prapro stands for “PRACE project”.

- Updates of the system generic information: architecture features, system software.
- Performance figures: R_{max}, R_{peak} (peak and achieved Linpack).
- System usage: number of PRACE Access projects, number of PRACE Access users, CPU hours used, job sizes, disk usage.
- User support: number of queries, number of queries of batch system.

2.2 System survey results

Both Tier-0 systems reported no specific updates for the generic information. The architecture features and system software/environments of JUGENE are well explained in [2] and the information for CURIE is provided in [3].

Table 1 shows the architecture, peak performance (R_{peak}), LINPACK performance (R_{max}) and the total number of cores on Tier-0 system JUGENE and CURIE. Note the R_{max} of CURIE is an estimated value. The definitive R_{max} of CURIE is expected to be available in June 2012.

Tier-0 Systems	*Architecture*	*R_{peak} (Gflop/s)*	*R_{max} (Gflop/s)*	*Cores*
JUGENE @ FZJ, Germany	IBM Blue Gene/P	1002700	825500	294912
CURIE @ TGCC, France	Bull system based on a modular x86 architecture with a mix of Fat Nodes (FN), Thin Nodes (TN), and Hybrid Nodes (HN)	FN: 104417 TN: 1700000 HN: 198161.6	FN: 87470 TN: > 1300000 HN: 109900	FN: 11520 TN: 10080 HN: 5040

Table 1 Tier-0 systems JUGENE and CURIE

Table 2 shows the number of PRACE Access projects and users on JUGENE and CURIE, including both the PRACE Preparatory Access and Regular Access, by the end date of survey periods, i.e. 31st December 2011 for JUGENE and 14th March 2012 for CURIE.

	JUGENE	*CURIE*
Total Preparatory Access projects	26	86
Total Preparatory Access users	47	238
Total Regular Access projects	39	22
Total Regular Access users	133	82

Table 2 Number of PRACE Access projects and users on Tier-0 systems

Table 3 lists the total CPU time usage for PRACE Preparatory Access, Regular Access and Industrial projects on JUGENE. The total numbers of jobs are also listed in the table. Within the survey period, PRACE Access consumed 35.3% of the total CPU core hours used on JUGENE, which is 95.3% of the total allocated CPU time for PRACE Access. The Regular Access projects consumed 98.8% of the total used CPU time for PRACE Accesses in the survey period, which is 34.8% of the total CPU time used on JUGENE.

On JUGENE	CPU Time Allocated (Core hours)	CPU Time Used (Core hours)	CPU Time Used out of Allocated (%)	CPU Time Used out of PRACE Total Used in Survey Period (%)	CPU Time Used out of JUGENE Total Used in Survey Period (%)	Total No. of Jobs
Total PRACE in Survey Period	1003670597	956866638	95.34%	NA	35.26%	71538
Preparatory Access	4050000	10105224	249.51%	1.06%	0.37%	4045
Regular Access	1404887264	944985842	67.26%	98.76%	34.82%	67236
Industry Projects	2000000	1775572	88.78%	0.19%	0.07%	257
Total on JUGENE in Survey Period	3271000597	2713634260	82.96%	NA	NA	659967

Table 3 CPU time usage and total number of jobs on JUGENE in the survey period

Figure 1 and Figure 2 show the details of CPU time allocated/used and the average CPU time cost per job for each project of Preparatory Access and Regular Access on JUGENE. More than half of the Preparatory Access projects on JUGENE used more than 100% of their allocated CPU core hours. 26 out of 39 Regular Access projects on JUGENE had used more than 50% allocated CPU time by the end date of survey period, giving an average of 67.3% CPU time used out of the total allocated time for PRACE Regular Access. For the finished Regular Access projects prapro01-prapro19, 101.6% of the total allocated CPU time has been used. The average CPU time per job varies quite widely between projects, ranging from a few thousand CPU hours to several hundred thousand.

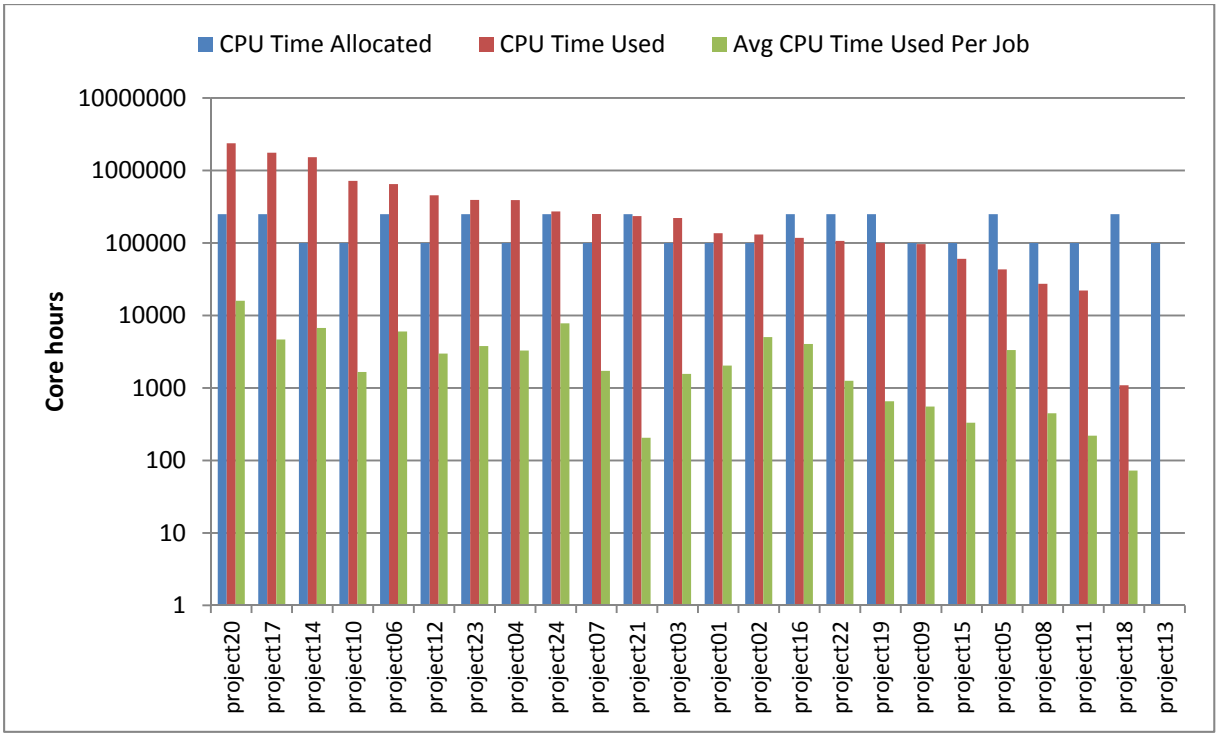

Figure 1 CPU time usage for Preparatory Access projects on JUGENE in the survey period

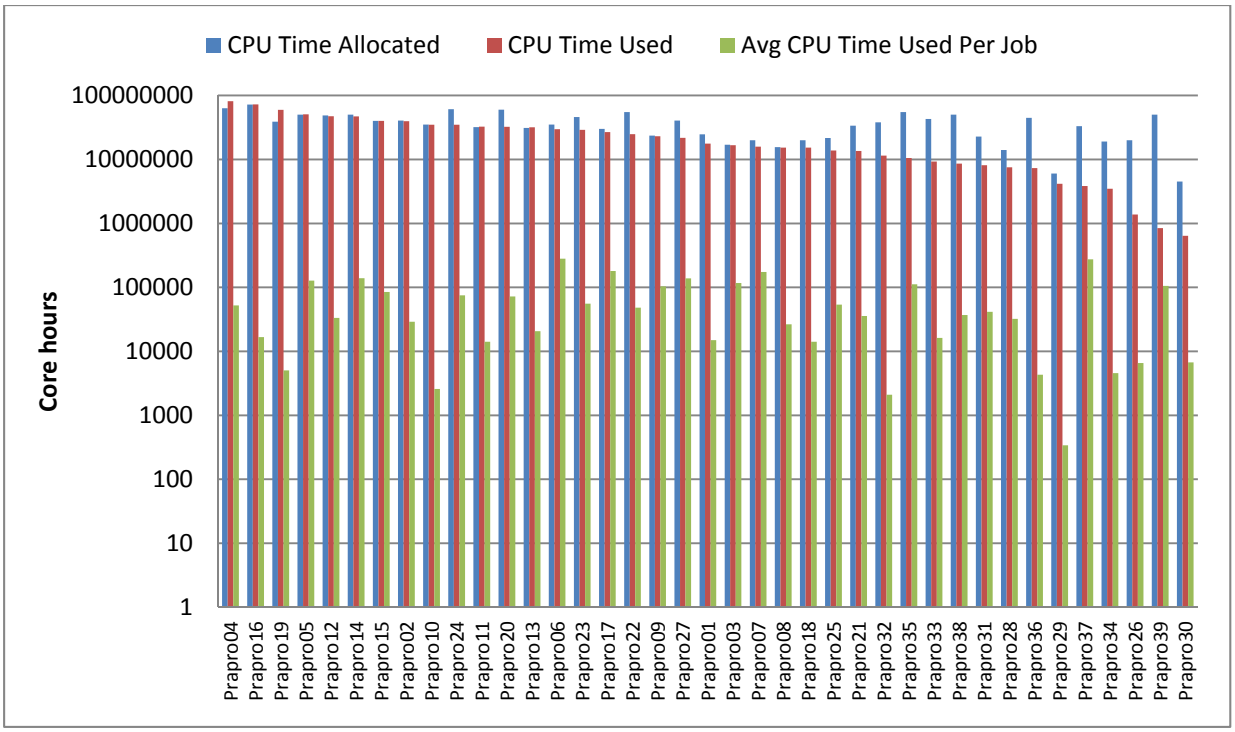

Figure 2 CPU time usage for Regular Access projects on JUGENE in the survey period

Table 4 lists the CPU time usage status and the total numbers of jobs for PRACE Preparatory Accesses and Regular Accesses on CURIE in the survey period of 1st March 2011 to 14th March 2012. The Preparatory Access projects had used 72.6% of allocated CPU time, which is around 21.5% out of the total PRACE CPU time usage on CURIE. The Regular Access projects had used only 12.4% of the allocated CPU core hours by the end date of survey period, which is around 78.5% out of the total PRACE CPU time usage on CURIE. However, it should be noted that all the Regular Access projects are still on-going during this survey time and will not finish until 1st November 2012. Around 65% of PRACE jobs were run by the Regular Access projects.

On CURIE	**CPU Time Allocated (Core hours)**	**CPU Time Used (Core hours)**	**CPU Time Used out of Allocated (%)**	**CPU Time Used out of PRACE Total Used in Survey Period (%)**	**Total No. of Jobs**
Total PRACE in Survey Period	254928699	38579504	15.13%	100%	78416
Preparatory Access	11432699	8299644	72.60%	21.51%	27252
Regular Access	243496000	30279860	12.44%	78.49%	51164

Table 4 CPU time cost and total number of jobs for PRACE Accesses on CURIE in the survey period

Figure 3 and Figure 4 shows the details of CPU time usage for each PRACE Access projects on CURIE. 22 Preparatory Access projects and 10 Regular Access projects had no CPU time used by the end date of survey period, and are therefore not shown on Figure 3 or Figure 4. For the 59 finished Preparatory Access projects, 103.0% of their allocated CPU core hours have been used. Again, there is a wide variation in the average number of CPU hours per job.

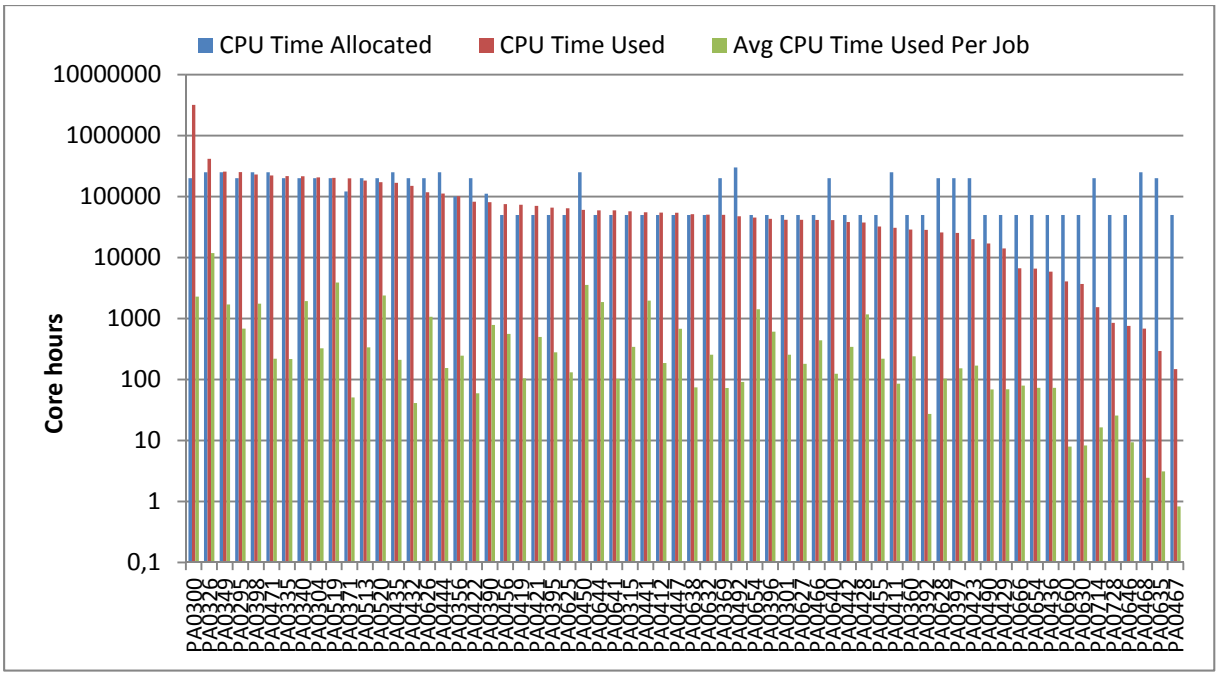

Figure 3 CPU time usage for Preparatory Access projects on CURIE in the survey period

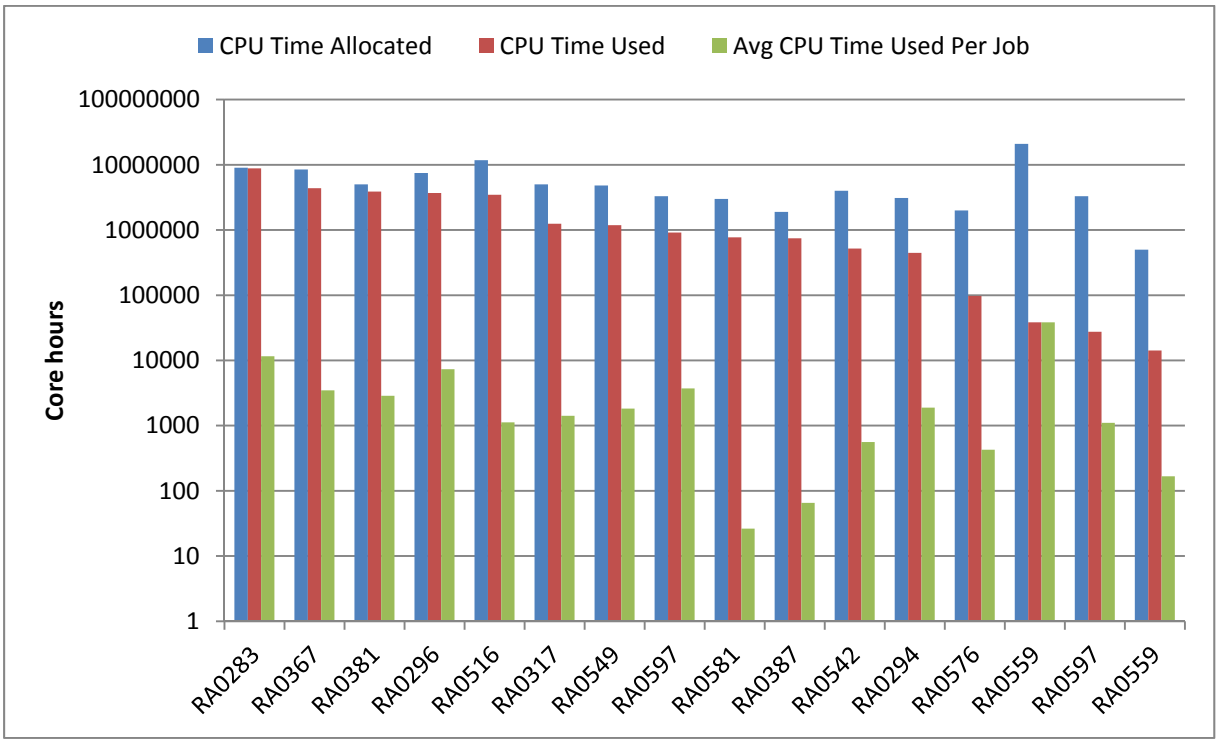

Figure 4 CPU time usage for Regular Access projects on CURIE in the survey period

Figure 5 shows the job size distribution of utilised cycles for each Regular Access project on JUGENE. The job sizes were divided in five ranges: up to 128 cores, 129-512 cores, 513-2048 cores, 2049-8192 cores and more than 8193 cores. Note that the distribution is expressed as a percentage of the utilised cycles, not as a percentage of the submitted jobs. 21 out of the 39 Regular Access projects on JUGENE spent more than 50% of their utilised cycles in the job size range of more than 8193 cores. 15 projects used more than 50% of utilised cycles in the job size range of 2049-9182 cores.

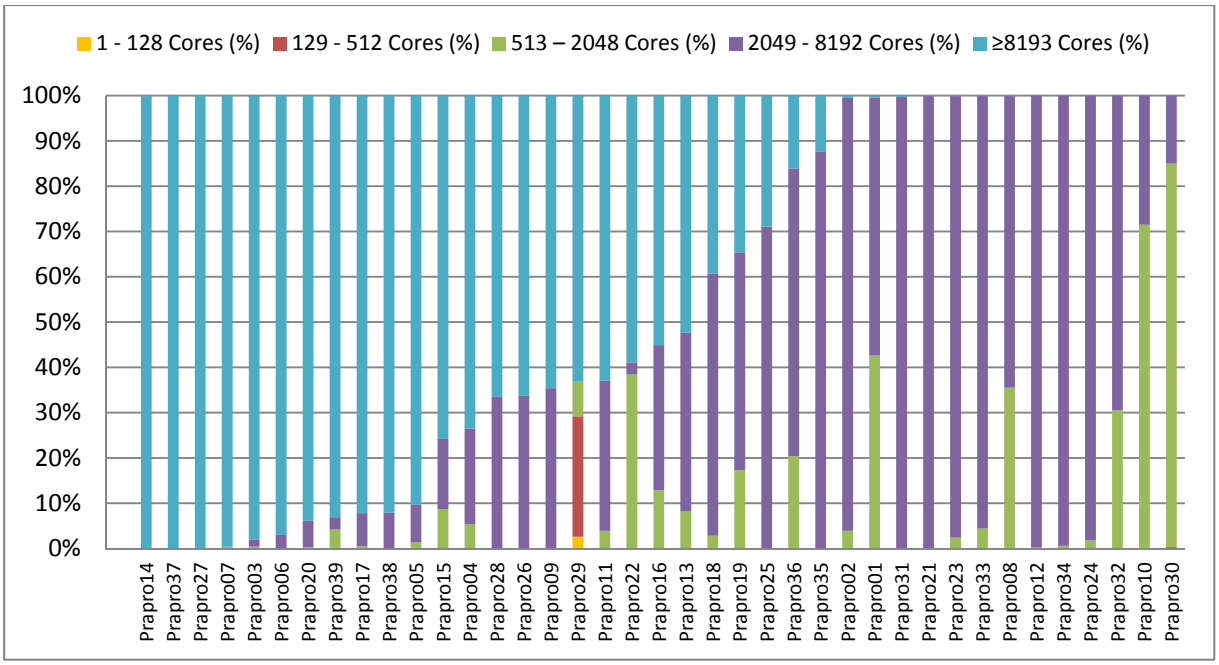

Figure 5 Job size distribution of utilised cycles for Regular Access projects on JUGENE

Figure 6 shows the aggregated distribution of utilised cycles in different job size ranges across all Regular Access projects on JUGENE. Just over 50% of the cycles were consumed by jobs running on more than 8192 cores. Around one third of the projects never used more than 8192 cores, but all made significant use of more than 2048 cores.

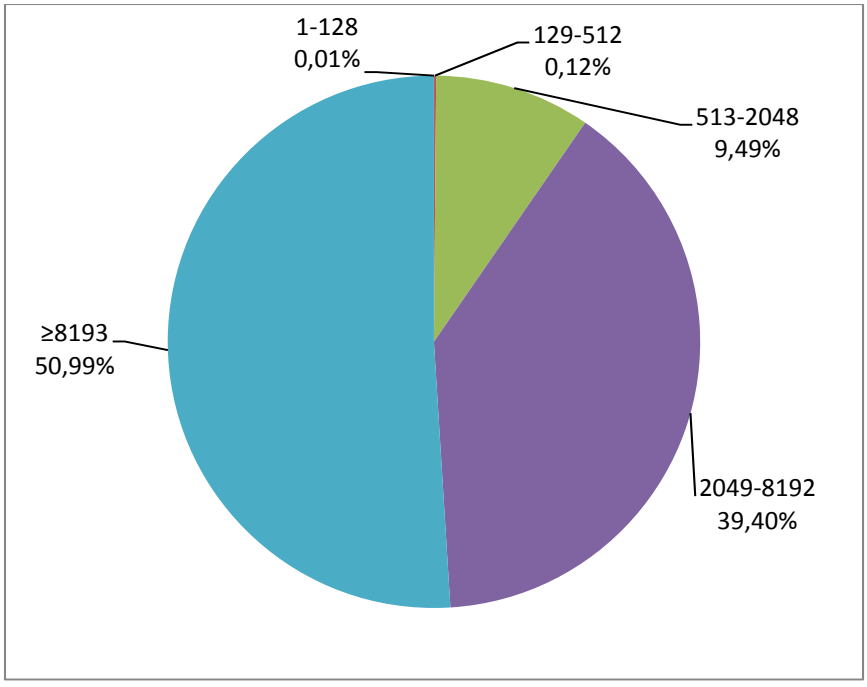

Figure 6 Aggregated distribution of utilised cycles by job size (Regular Access) on JUGENE

Figure 7 shows the job size distribution of utilised cycles for each Regular Access project on CURIE in the five ranges of job size. 10 out of the 16 projects used more than 50% of utilised cycles in the job size range of 513-2048 cores, and just two projects ran jobs with only 512 or fewer cores. Figure 8 shows the aggregated job size distribution across all Regular Access projects on CURIE. 45% of the cycles used were in jobs of over 2048 cores.

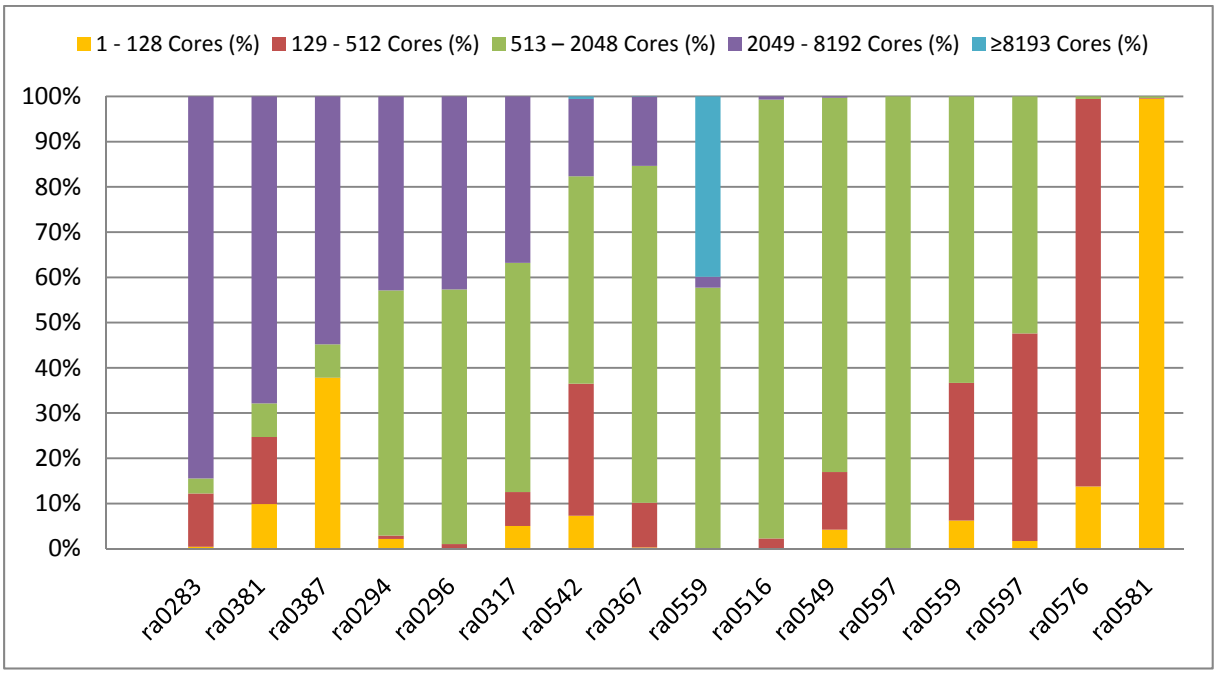

Figure 7 Job size distribution of utilised cycles for Regular Access projects on CURIE

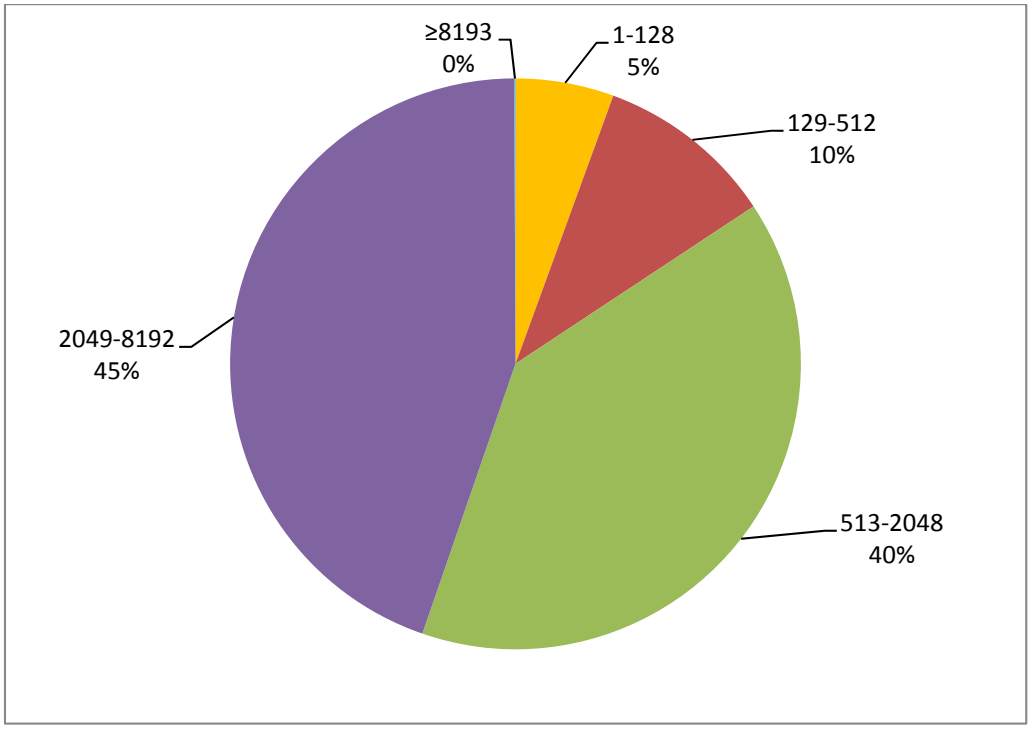

Figure 8 Aggregated distribution of utilised cycles by job size (Regular Access) on CURIE

A wide range of disk usage behaviour was observed on JUGENE and CURIE. Figure 9, Figure 10 and Figure 11 are the Regular Access disk usage status of Home directory, Work directory and Archive on JUGENE at the end date of survey period, i.e. 31st December 2011. Of course, these figures only represent a snapshot in time of disk usage.

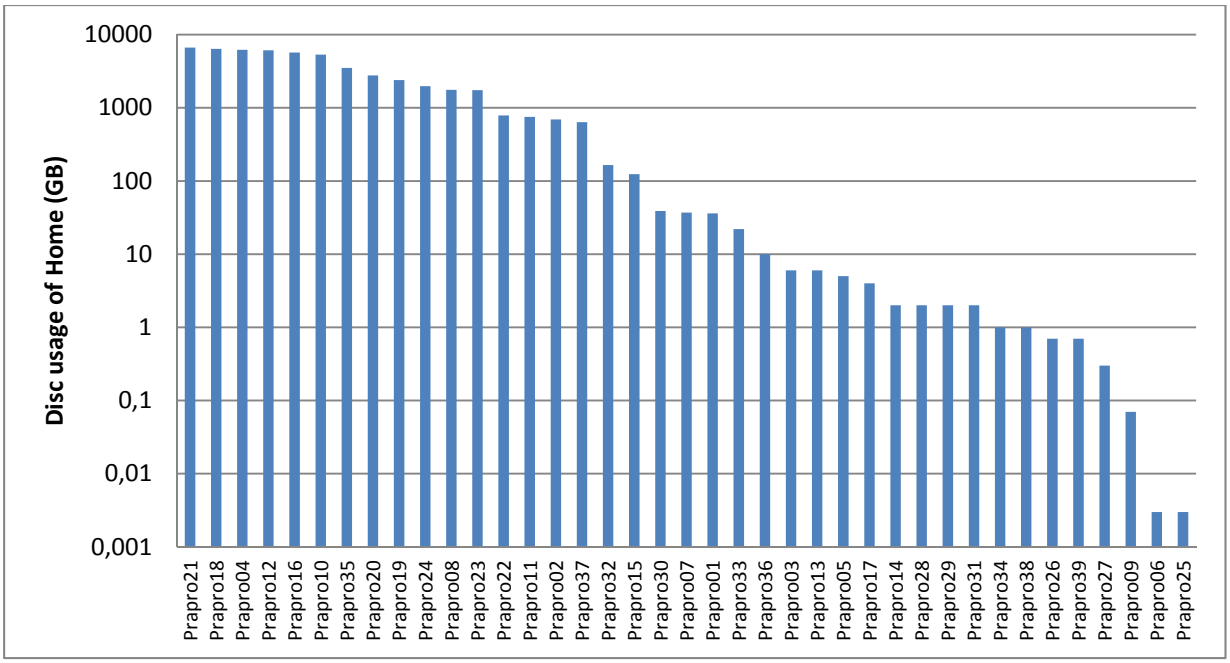

Figure 9 Disk usage of Home directory for Regular Access on JUGENE

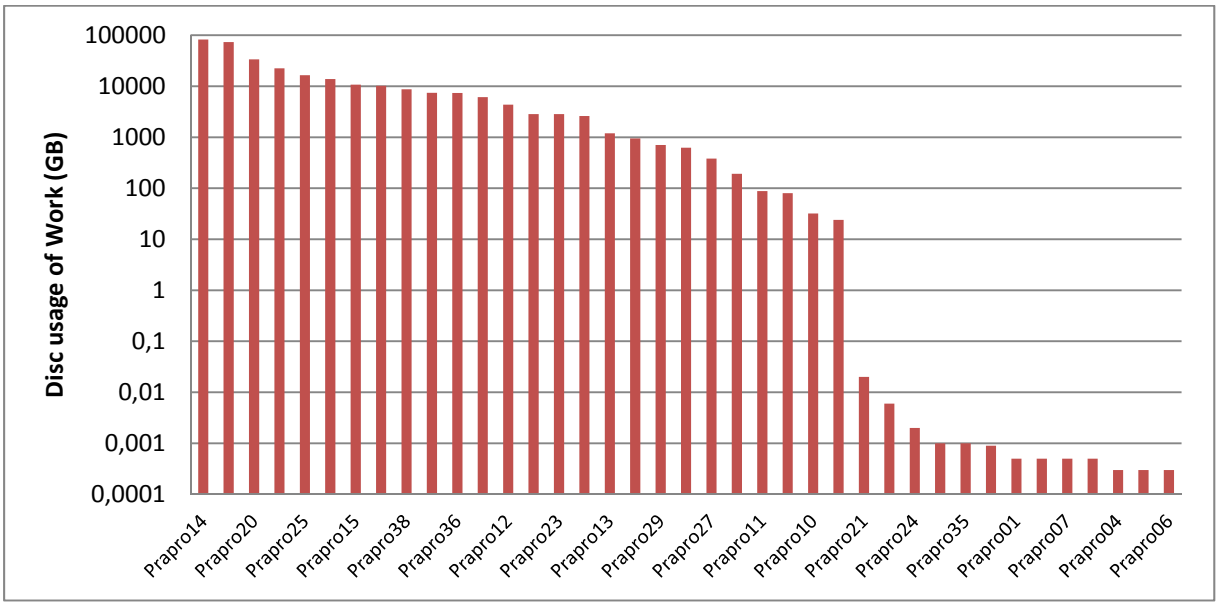

Figure 10 Disk usage of Work directory for Regular Access on JUGENE

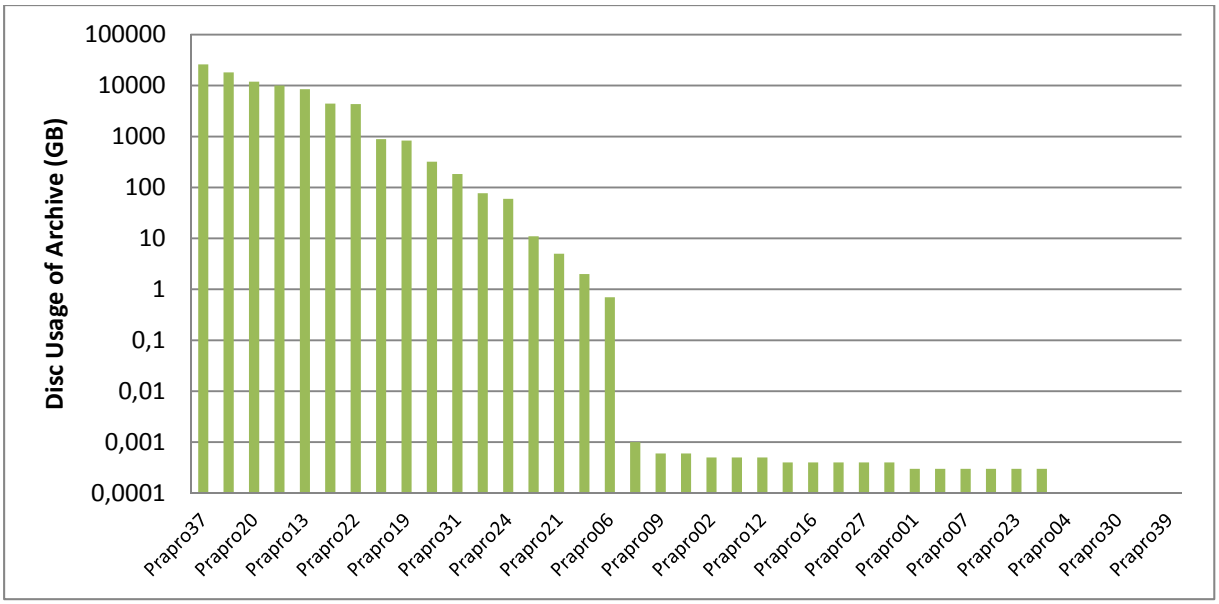

Figure 11 Disk usage of Archive for Regular Access on JUGENE

Figure 12, Figure 13, Figure 14 and Figure 15 are the Regular Access disk usage status of HOME directory, SCRATCH directory, WORK directory and STORE on CURIE at the end date of survey period, i.e. 14th March 2012.

On both systems a wide range of disk usage is observed, and a small number of projects typically account for most of the disk usage.

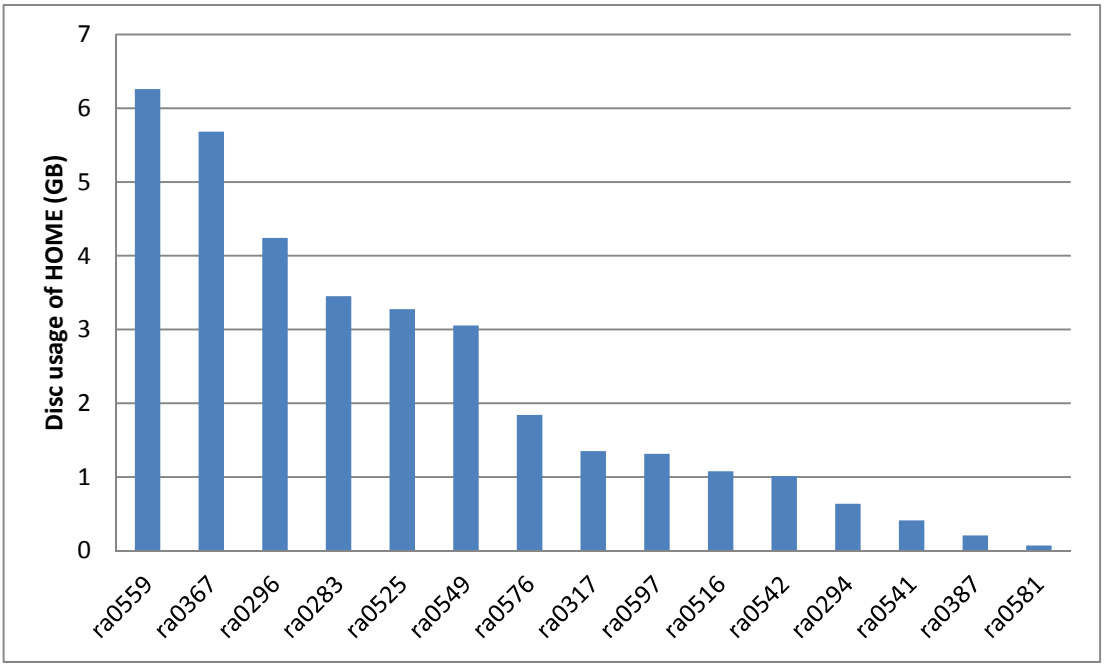

Figure 12 Disk usage of HOME for Regular Access on CURIE

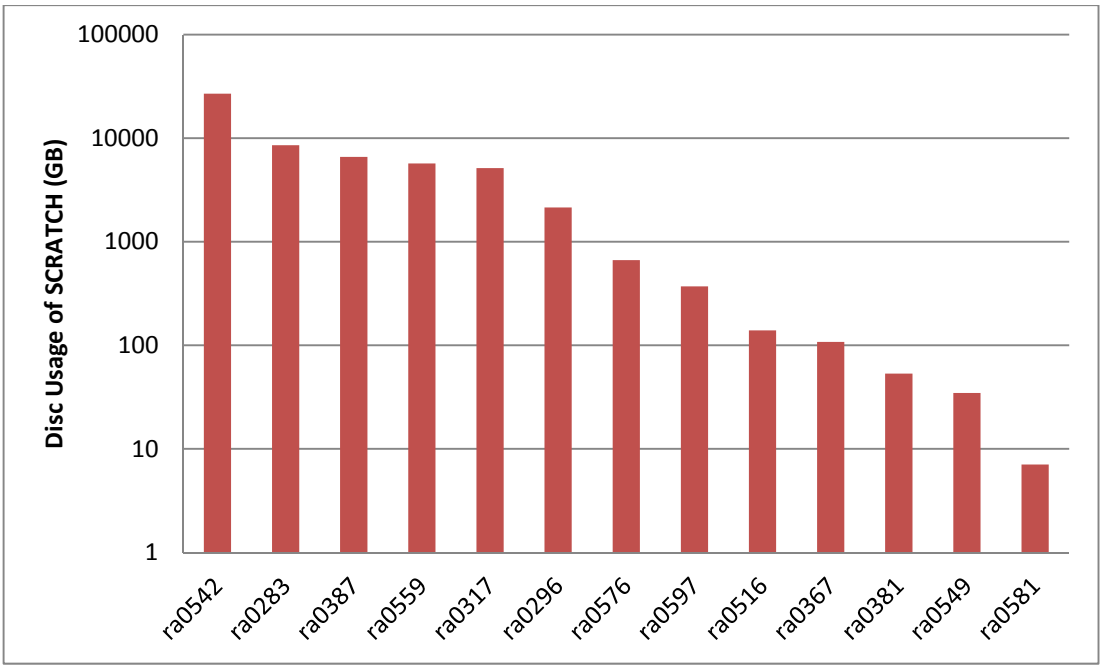

Figure 13 Disk usage of SCRATCH for Regular Access on CURIE

Figure 14 Disk usage of WORK for Regular Access on CURIE

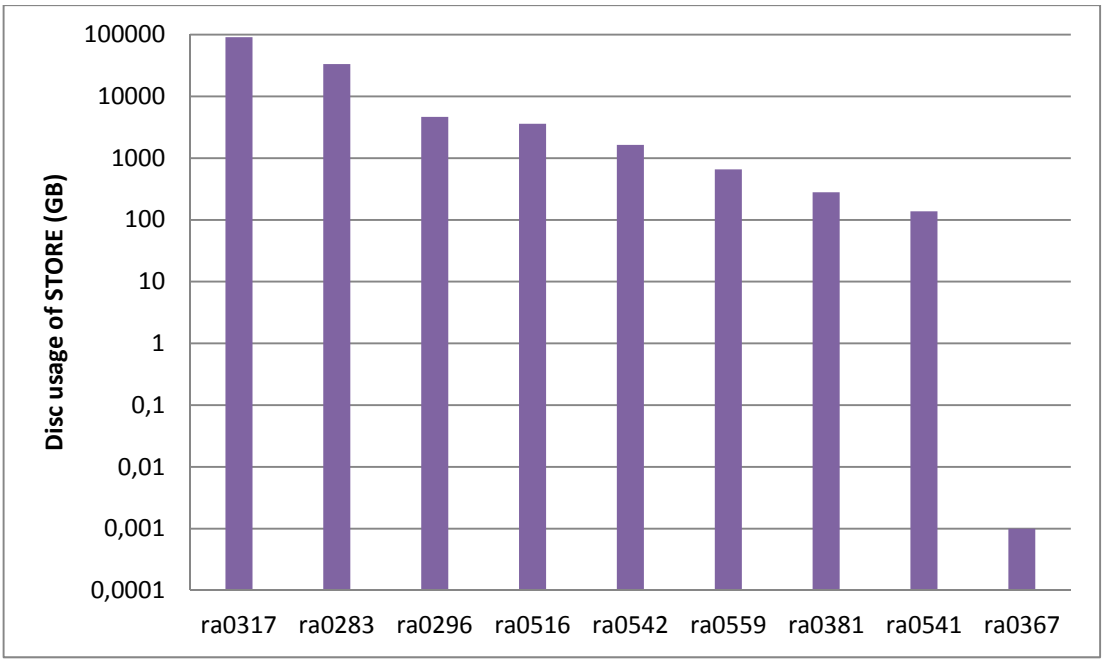

Figure 15 Disk usage of STORE for Regular Access on CURIE

With respect to user support, there were 673 queries received from PRACE Preparatory and Regular Access users on CURIE of which 20% were related to the batch system/scheduler. User support query statistics were not available for JUGENE.

2.3 Summary

We have surveyed the PRACE Access usage of the two Tier-0 systems JUGENE and CURIE. The principal findings from this survey are:

- PRACE usage accounts for just over one third of all the CPU hours consumed on JUGENE. This matches the commitment of one third of JUGENE capacity to PRACE.
- Completed Regular Access projects have used all their allocated CPU hours.
- Just over 50% of the cycles used by Regular Access projects on JUGENE were in jobs running on more than 8192 cores.
- 45% of the cycles used by Regular Access projects on CURIE were in jobs running on more than 2048 cores.
- A small number of projects account for a large fraction of the disk usage on both systems.
- Applications exhibit a wide range of disk requirements for running production jobs.

3 User Survey

3.1 Overview

The user survey was designed to understand the users' experience and requirements of the PRACE Tier-0 system usage. The survey consisted of a set of 48 questions that were to be answered by the users of PRACE Tier-0 systems who were, or had been, involved in the PRACE Access projects, including both Preparatory Access and Regular Access projects. All such users were invited, via emails sent by the Tier-0 hosting sites, to respond to the survey. The responses were collected between 13th December 2011 and 5th March 2012. A total of 62 valid responses were received. Note that some projects were therefore not represented in the survey, and in a few cases, more than one user from a project responded to the survey.

The survey was implemented using SurveyMonkey (www.surveymonkey.com) and the results were obtained in the form of Excel spreadsheets. These spreadsheets are available for use by other work packages/tasks in PRACE-1IP as required.

3.2 User survey results

The responses to each question are summarised below. Information includes:

- The response rate, i.e. the fraction of respondents who answered this particular question;
- A summary of the valid responses. Figures/Tables are given where possible. For the single-choice questions, a pie chart will usually be given. For the multi-choice questions, a column chart will usually be given.
- Comments on the results, where appropriate.

Question 1: Your organisation &
Question 2: Your country of residence (Single choice) &
Question3: Your email (optional, confidential)

Response rate: 62/62 to Question 1; 62/62 to Question 2; 29/62 to Question 3

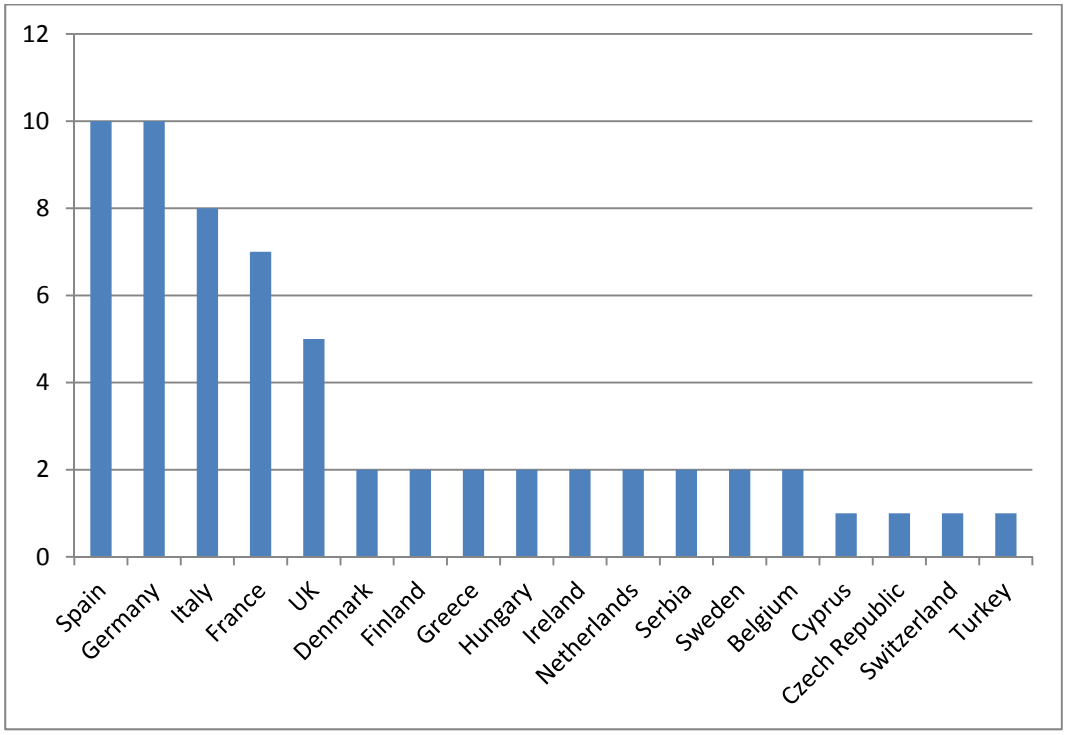

Figure 16 Summary of responses to Question 2

Comment: For confidentiality reasons, the responses were summarised by the country of residence only.

Question 4: Which system(s) are / were you working on for your PRACE project(s)?

Response rate: 62/62

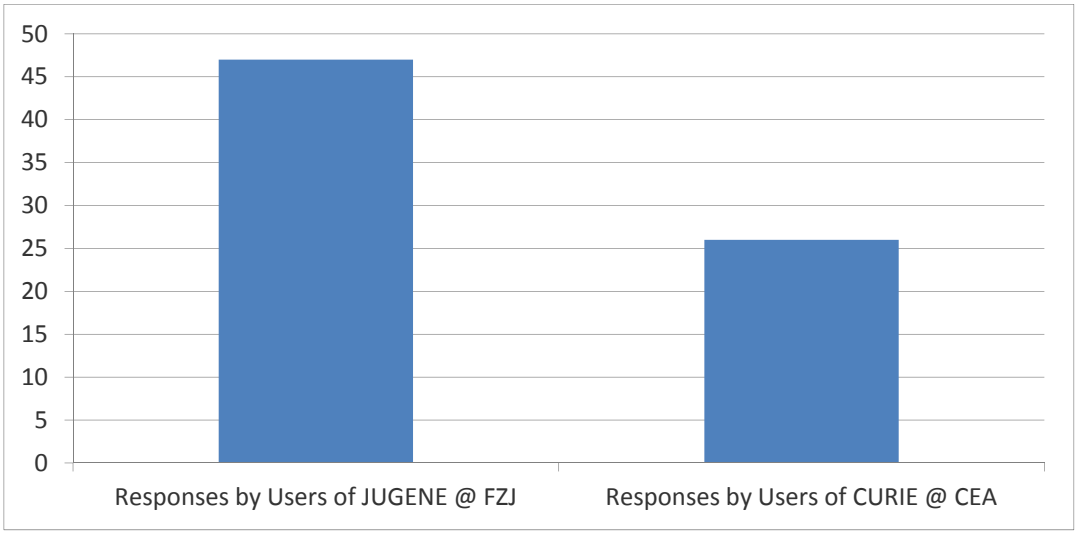

Figure 17 Summary of responses to Question 4

Comment: Note some users have accounts on both JUGENE and CURIE.

Question 5: Your PRACE project type? (Single choice)

Response rate: 62/62

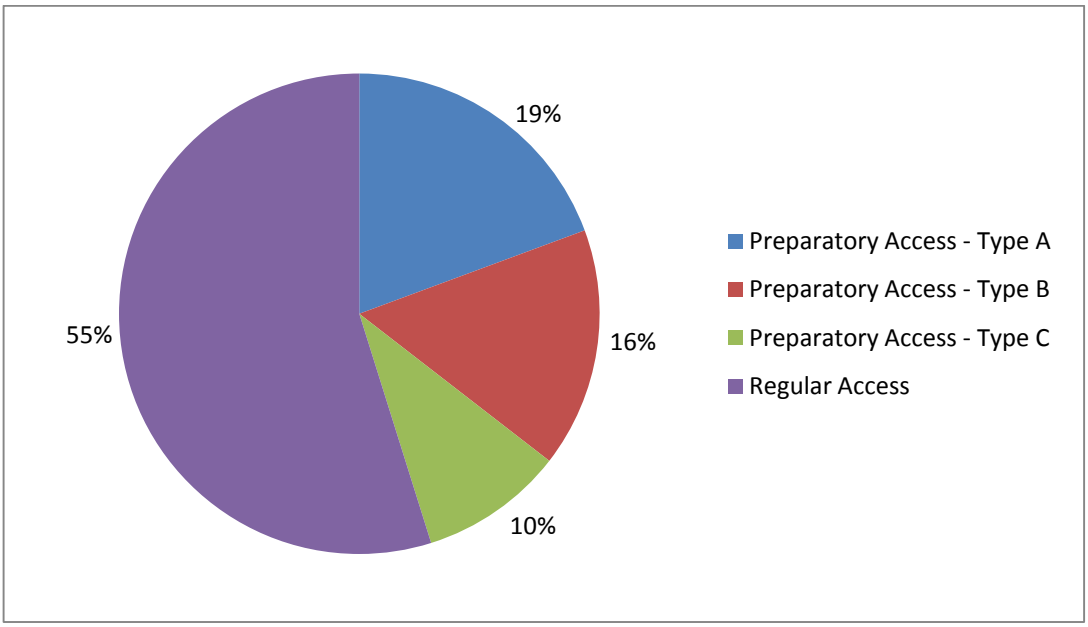

Figure 18 Summary of responses to Question 5

Question 6: If your PRACE project type is Preparatory Access Type A, please select your project name from the following dropdown list.

Response rate: 9/12 respondents who selected Preparatory Access Type A in Question 5, answered this question.

PRACE Preparatory Access Project (Type A)	Response number
Structure and evolution of an active region on the Sun	1
NadiaSpectral	2
The Genetic Hybrid Algorithm	1
petascaing of High Resolution EC-EARTH	1
PRACE Project 2010PA0630 on CURIE	1
Developing improved models for Tsunamis	1
Performance and Scalability of HadGEM2-ES and HadGEM2 Configurations of the UK Met. Office Unified Model	1
Insights into the mechanism of oncogenesis of the mutant protein PI3Ka from Molecular Dynamics simulations	1
Total responses	*9*
Number of Preparatory Access projects of Type A in responses	*8*
Total number of Preparatory Access projects of Type A	*23*

Table 5 Summary of responses to Question 6

Comment: There were total 23 Preparatory Access projects of Type A by the end of this survey. A few users indicated they were involved in the Preparatory Access projects of Type A in the previous Question 5 but gave no project name in this question.

Question 7: If your PRACE project type is Preparatory Access Type B, please select your project name from the following dropdown list.

Response rate: 9/10 respondents who selected Preparatory Access Type B in Question 5, answered this question.

PRACE Preparatory Access Project (Type B)	Response Number
Climate, Organized and Unorganized Atmospheric Convection	2
noFUDGE: Flow Unsteadiness computed by DG finite Elements	2
Three-Dimensional simulation of high frequency wave propagation in a Mode Converter (3DMC)	1
Aero-Elastics with Distributed Octrees	1
MD in electromagnetic fields	1
Non-adiabatic molecular dynamics with explicitly treated electronic degrees of freedom.	1
Scale-out of the Propag Electrocardiology Code on Petascale Architectures	1
Total responses	*9*
Number of Preparatory Access projects of Type B in responses	*7*
Total number of Preparatory Access projects of Type B	*16*

Table 6 Summary of responses to Question 7

Comment: There were total 16 Preparatory Access projects of Type B by the end of this survey. One user who selected the Preparatory Access projects of Type B in the previous Question 5 didn't give the project name.

Question 8: If your PRACE project type is Preparatory Access Type C, please select your project name from the following dropdown list.

Response rate: 6/6 respondents who selected Preparatory Access Type C in Question 5, answered this question.

PRACE Preparatory Access Project (Type C)	Response Number
Self organization, pattern formation and morphological instabilities in suspensions of microswimmers	2
Optimizing a 6D global Vlasov simulation of Earth's magnetosphere	3
CP2K (Internal projects 7.1 and 7.2- 1IP)	1
Total responses	*6*
Number of Preparatory Access projects of Type C in responses	*3*
Total number of Preparatory Access projects of Type C	*12*

Table 7 Summary of responses to Question 8

Comment: Table 7 lists all the Preparatory Access projects of Type C in users' responses. There were total 12 Preparatory Access projects of Type C by the end of this survey.

Question 9: If your PRACE project type is Regular Access, please select your project name from the following dropdown list.

Response rate: 26/34 respondents who selected Regular Access in Question 5, answered this Question.

PRACE Regular Access Project	Response Number
Ab initio molecular dynamics simulations of proton transport in a biological ion channel	1
Entrainment effects in rough-wall boundary layers	1
QCD Thermodynanics with Wilson fermions	1
Investigating the effects of quantum nuclear motion in an enzyme that employs hydrogen tunnelling	1
Turbulent entrainment due to a plume impinging on a density interface	2
Non diffusive transport in ITG plasma turbulence	1
Large scale high resolution blood flow simulations in realistic vessel geometries	1
Ab Initio Modeling of Solar Active Regions	1
REFIT - Rotation effects on flow instabilities and turbulence	2
Large-Eddy Simulation of high-frequency instabilities under transcritical conditions	1
Large Scale simulations of Ly-alpha and Ly-break galaxies in the high-z universe: Probing the epoch of reionization	1
Diversity of Type Ia supernovae from initial conditions of the exploding white dwarf star	2
Structural and conformational requisites in the folding process of the DNA quadruplexaptamer TBA	1
Extreme Earthquake Wave Propagation Modelling (E2WPM)	2
The molecular bases of the transport cycle of APC antiporters	1
MS-COMB: Multi-Scale Analysis and Numerical Strategies for the Simulation of Premixed Turbulent Combustion in Realistic Geometries	1
Pushing the Strong Interaction past its Breaking Point: QCD in the quark-gluon plasma phase.	1

Thermal Dilepton Rates and Electrical Conductivity in the Quark Gluon Plasma	1
Structure of turbulence in supersonic boundary layers at high Reynolds number	1
Protein effects on the structural and optical properties of biological chromophores: Quantum Monte Carlo / Molecular Mechanics calculations on Rhodopsin and Light Harvesting Complexes.	1
Entrainment effects in rough-wall boundary layers	1
Pushing the Strong Interaction past its Breaking Point	1
Total responses	*26*
Number of Regular Access projects in responses	*22*
Total number of Regular Access projects	*26*

Table 8 Summary of responses to Question 9

Comment: There were total 26 Regular Access projects by the end of this survey. A few users indicated their project type as Regular Access in the previous Question 5 but provided no project name.

Question 10: Which project account was used on each system for your project?

Response rate: 48/62

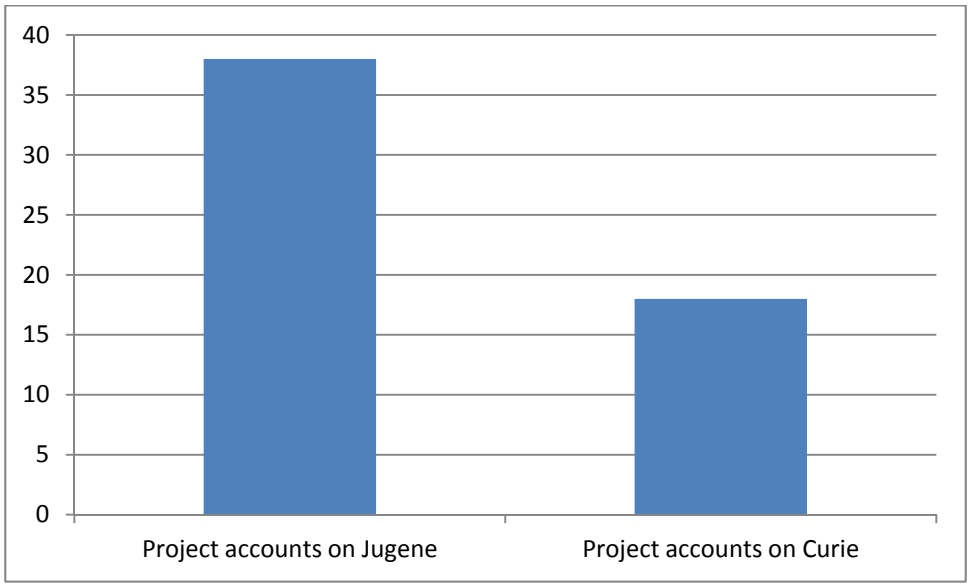

Figure 19 Summary of responses to Question 10

Comment: A few responses provided project accounts on both JUGENE and CURIE.

Question 11: Which individual account (user ID) was used on each system?

Response rate: 42/62

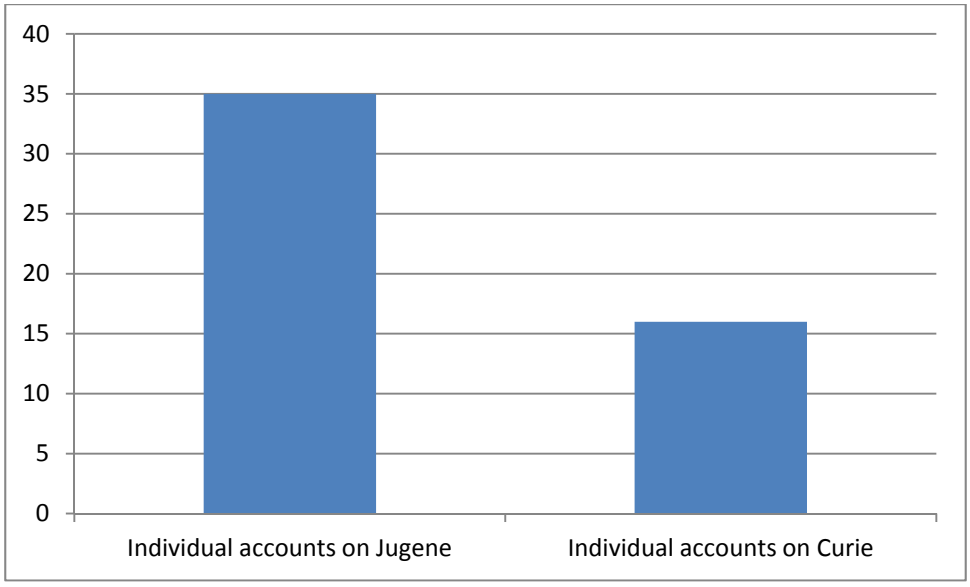

Figure 20 Summary of responses to Question 11

Comment: A few of responses provided individual accounts on both JUGENE and CURIE. Some responses provided multiple individual accounts on one system.

Question 12: What is your role in the project? (Multiple choice)

Response rate: 48/62

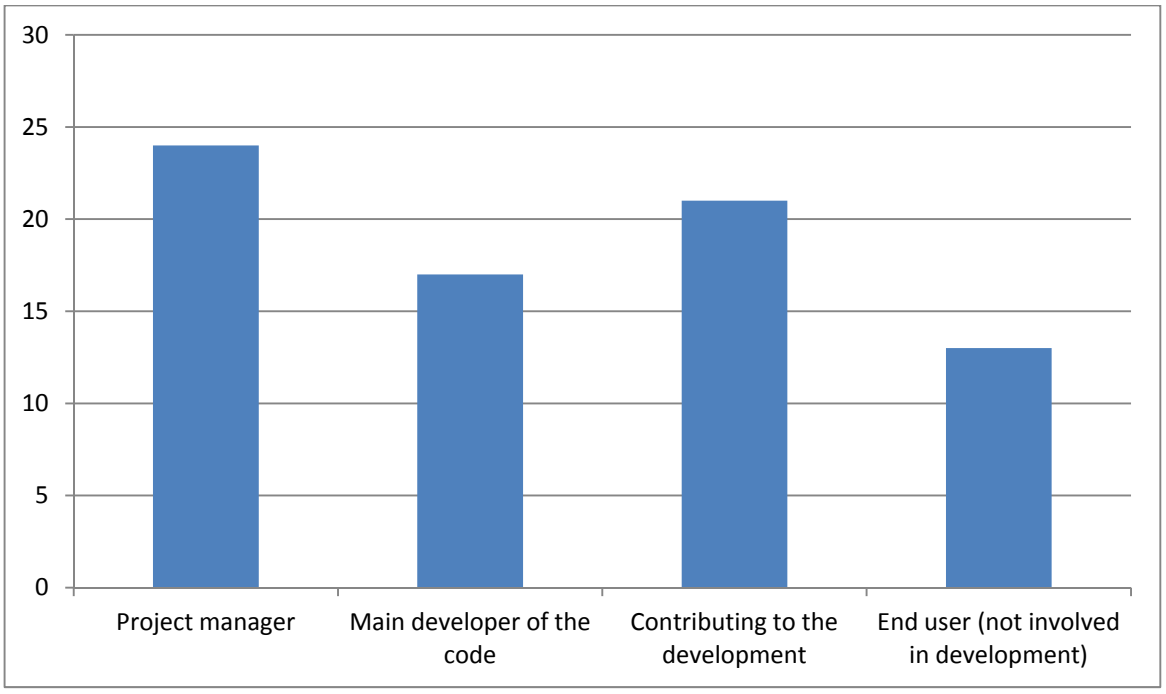

Figure 21 Summary of responses to Question 12

Comment: The responses to this question show that all four roles are well represented by the users completing the survey.

Question 13: What is the name of the application code used in your PRACE project?

Response rate: 54/62

Project Type	Applications Name	Response Count
Preparatory Access Type A	Alya	1
Preparatory Access Type B	Apes	1
Preparatory Access Type B	Argo	2
Preparatory Access Type C	CP2K	1
Preparatory Access Type A	EC-EARTH	1
Preparatory Access Type A	GHA	1
Preparatory Access Type B	HALO3D	1
Preparatory Access Type B	IPSL-CM	2
Preparatory Access Type B	LAMMPS	1
Preparatory Access Type C	Ludwig	2
Preparatory Access Type B	MULTIDIS	1
Preparatory Access Type A	nadiaADI or nadiaCN	1
Preparatory Access Type A	NadiaSpectral	1
Preparatory Access Type A	NAMD	1
Preparatory Access Type A	Pencil Code	1
Preparatory Access Type B	Propag	1
Preparatory Access Type A	Tsunamis	1
Preparatory Access Type A	Unified Model (UK Met Office)	1
Preparatory Access Type C	Vlasiator	2
Regular Access	AHF	1
Regular Access	ARTIS	1
Regular Access	Chroma	3
Regular Access	CPMD	2
Regular Access	DD-HMC	1
Regular Access	Dynqcd	1
Regular Access	E2WPM	1
Regular Access	EUTERPE	1
Regular Access	fpDNS_AW	1
Regular Access	Gadget 3	1
Regular Access	Ginnungagap	1
Regular Access	Gromacs	2
Regular Access	LB3D	2
Regular Access	LEAFS	1
Regular Access	MP-HMC	1
Regular Access	NAMD	1
Regular Access	PhotonPlasma	1
Regular Access	Simson	1
Regular Access	Stagger Code	1
Regular Access	TBL	2
Regular Access	TurboRVB	1
Regular Access	YALES2	1
Total number of application names		*40*
Total responses		*54*

Table 9 Summary of responses to Question 13

Comment: Note that NAMD was used in one Preparatory Access project of Type A and in one Regular Access project as well, so it is listed twice in the table above. A few users gave no name in the responses but did provided information on scientific areas, parallelisation methods and programming languages for the following Questions (14, 15 and 16).

Question 14: Please select one of the following scientific areas that your application belongs to. (Single choice)

Response rate: 54/62

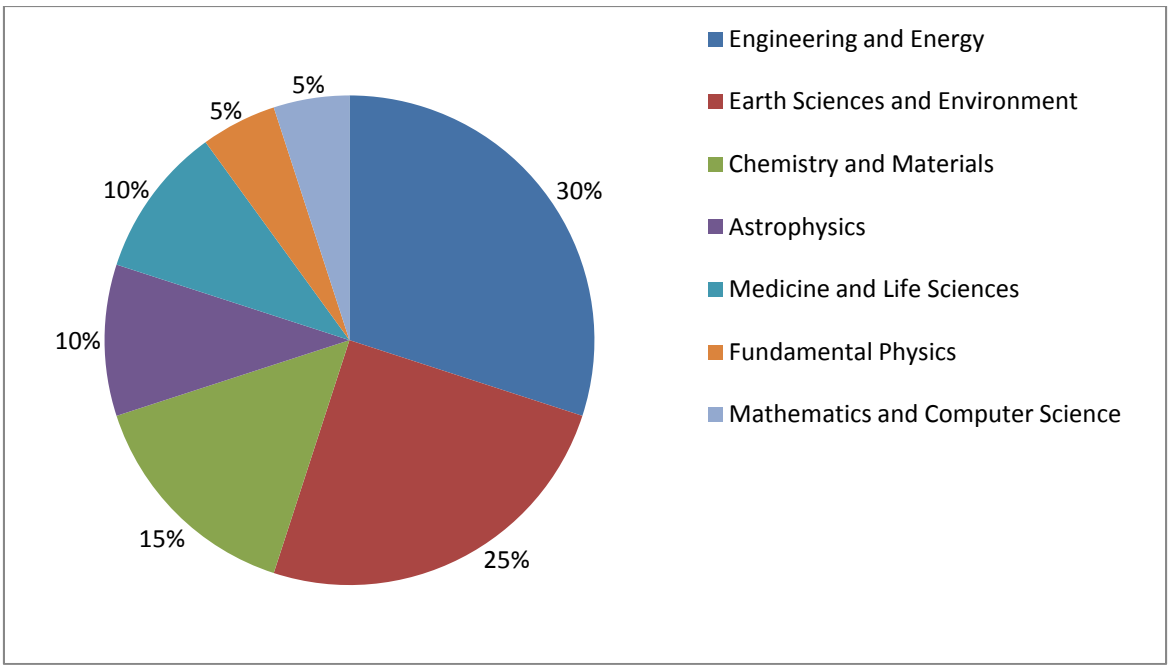

Figure 22 Summary of responses to Question 14 for Preparatory Access

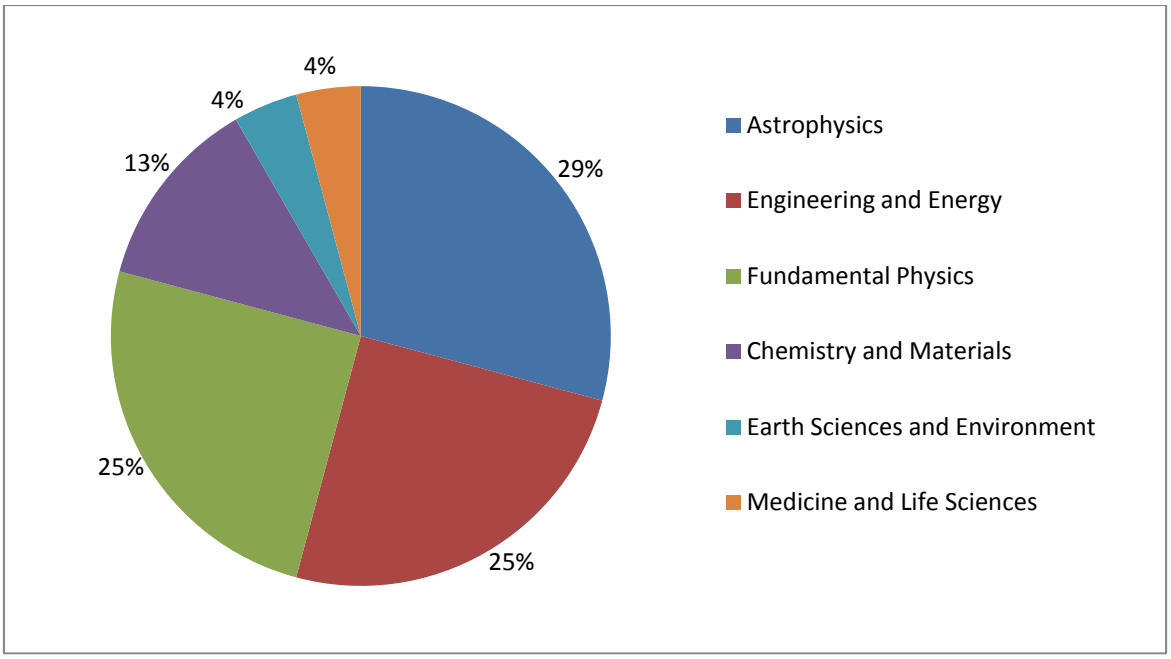

Figure 23 Summary of responses to Question 14 for Regular Access

Comment: Figure 19 and Figure 20 show the distribution of application scientific areas in Preparatory Access projects and Regular Access projects. Note the distribution is shown as a percentage of the applications rather than as a percentage the number of survey responses. There was one response of Preparatory Access and two responses of Regular Access providing no application name. It was assumed that they are not duplicated with other applications in the pie charts above.

Question 15: Which parallelisation method does your application use? (Single choice)

Response rate: 54/62

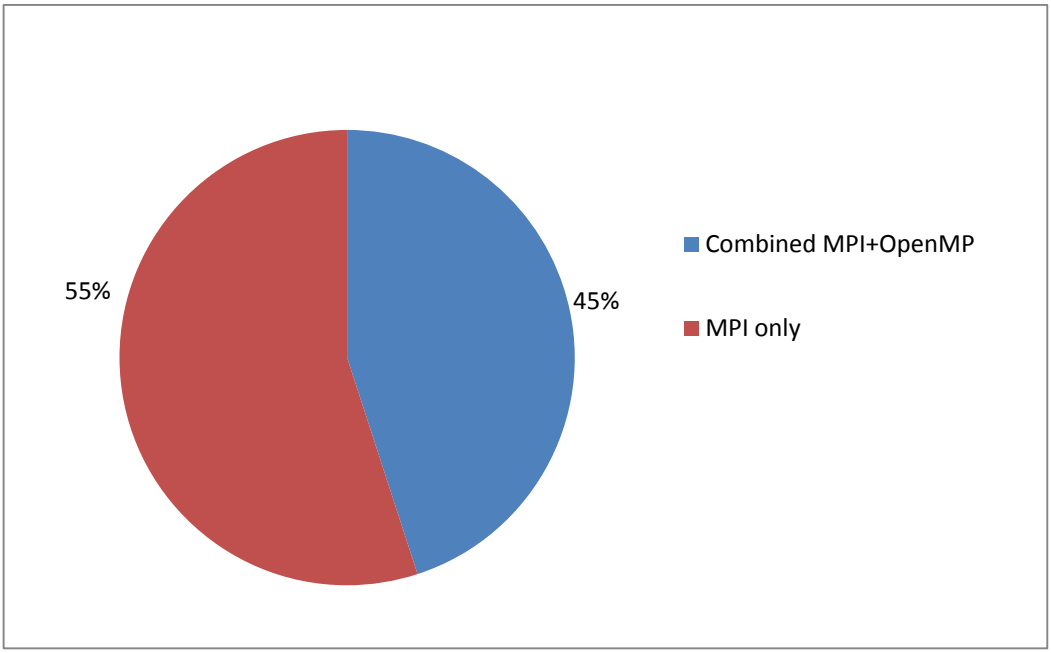

Figure 24 Summary of responses to Question 15 for Preparatory Access

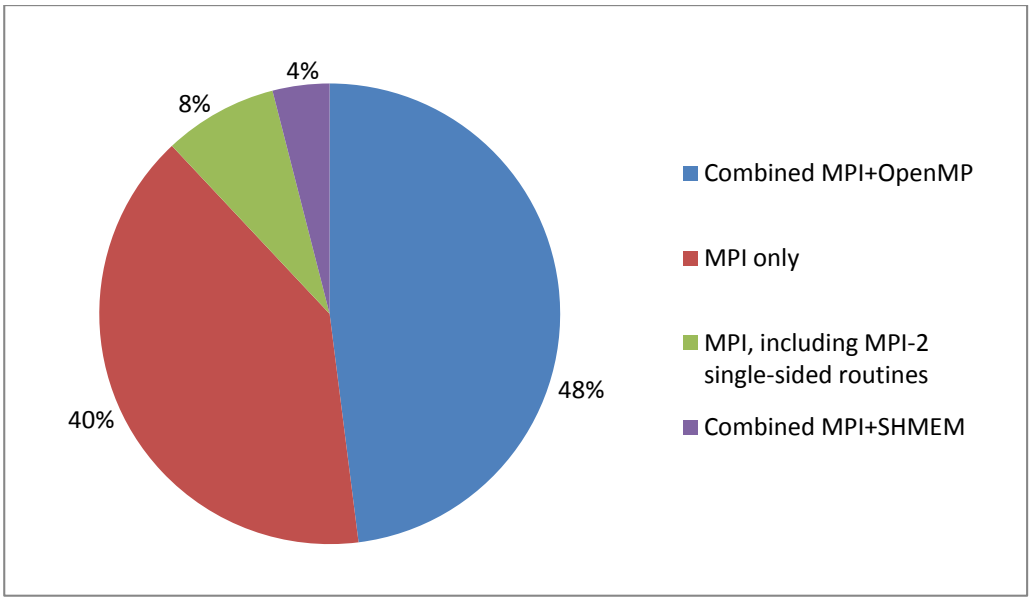

Figure 25 Summary of responses to Question 15 for Regular Access

Comment: Figure 21 and Figure 22 show the distribution of application parallelisation methods in Preparatory Access projects and Regular Access projects. Note the distribution is based on the applications rather than the number of survey responses. Parallelisation using combined MPI+OpenMP or MPI only takes the major percentage of all the parallel implementations. The proportion of applications using combined MPI+OpenMP reported here is much higher than in previous surveys carried out by the PRACE project [4][5]. This may reflect the importance of this combined model in permitting high scalability on Tier-0 systems. No applications were reported as using other parallelization models (e.g. PGAS).

Question 16: Which programming languages does your application use? (Multiple choice)

Response rate: 54/62

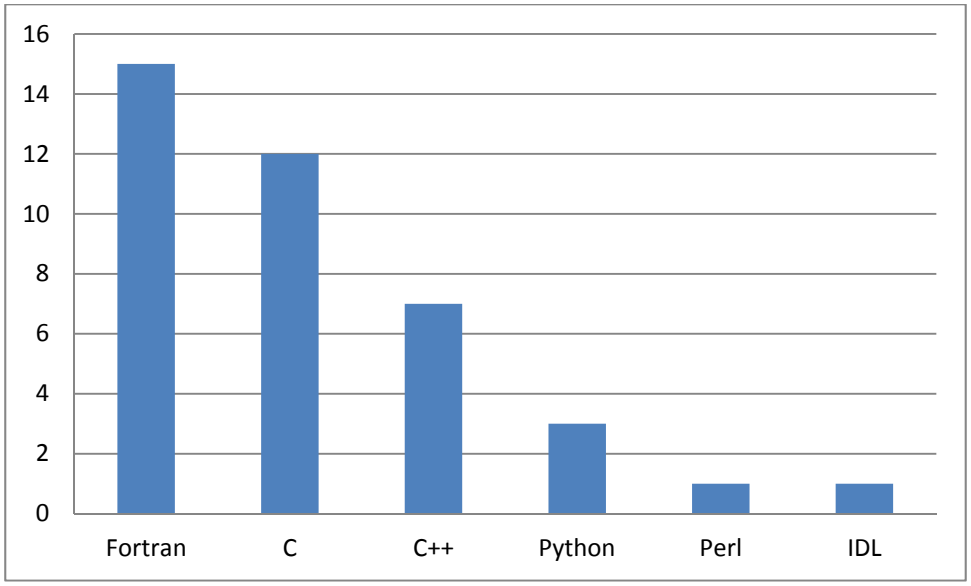

Figure 26 Summary of responses to Question 16 for Preparatory Access

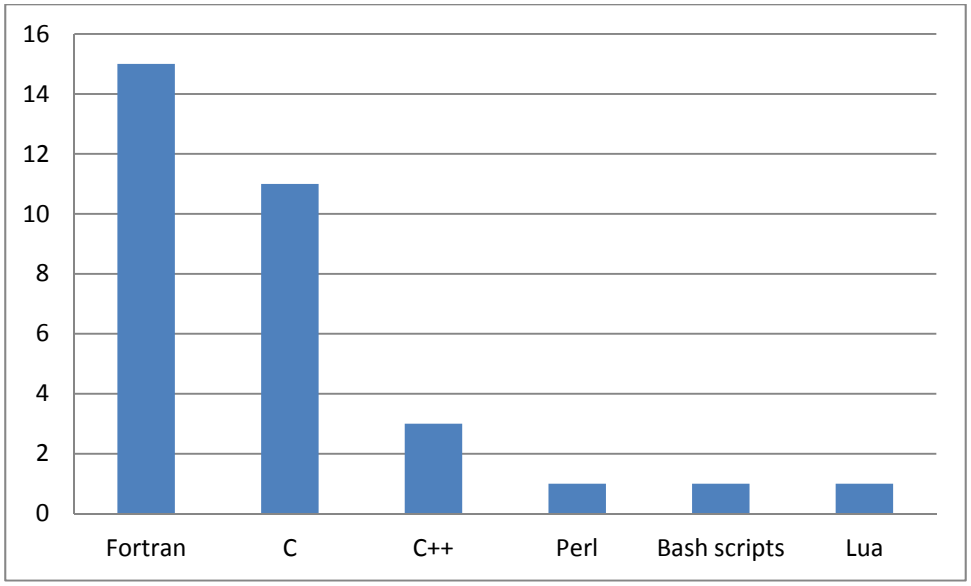

Figure 27 Summary of responses to Question 16 for Regular Access

Comment: Fortran and C/C++ are still the most popular programming languages used in developing the applications, with a roughly equal split between Fortran and C/C++.

Question 17: What memory size per core is required for your typical production jobs (in GB)?

Response rate: 34/62

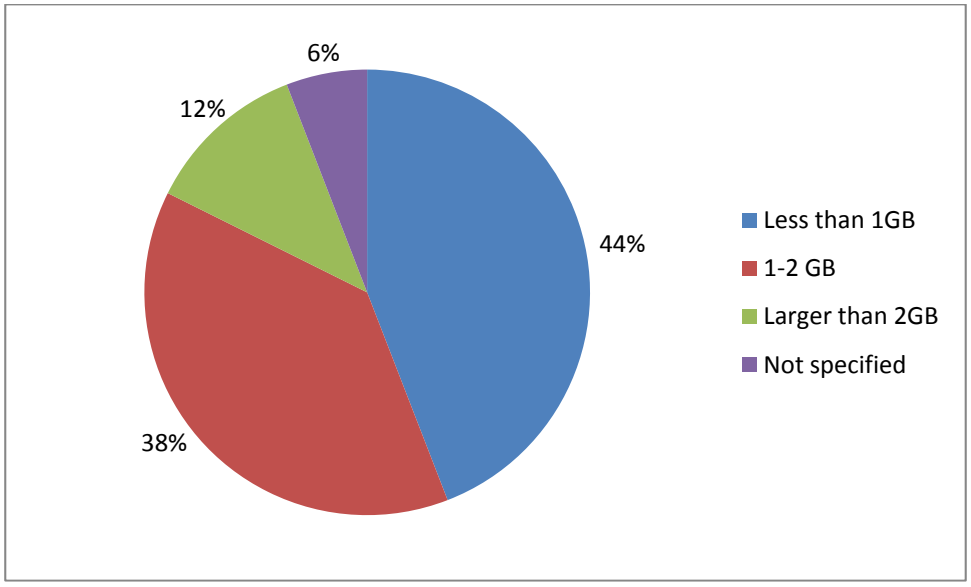

Figure 28 Summary of responses to Question 17

Comment: Around 70% of the responses to this question indicated that the memory size per core required for their production jobs is less than 1GB. However, this may be influenced by the fact that jobs running on JUGENE only have access to 0.5 GB per core in any case.

Question 18: What is the minimum amount of disk space required per production job (in GB)?

Response rate: 34/62

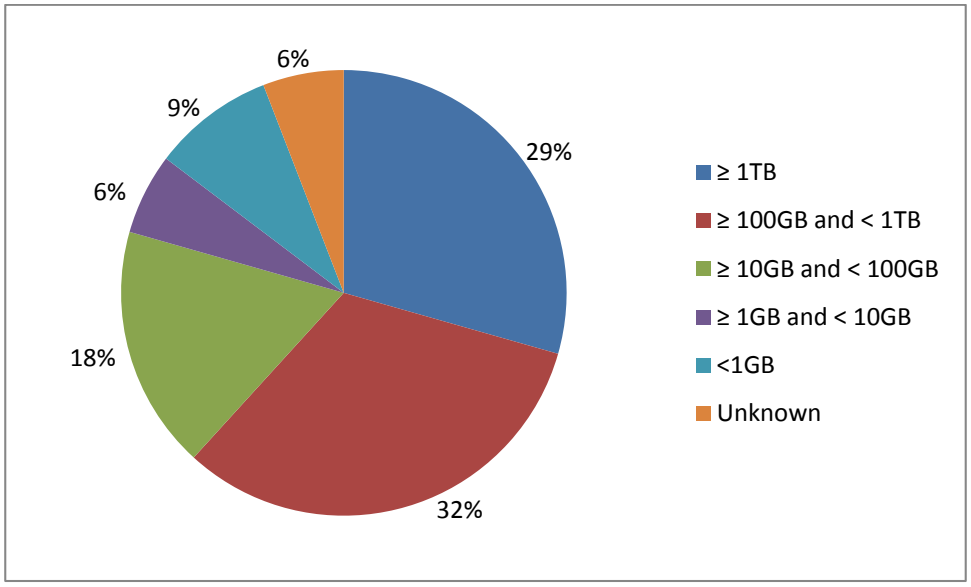

Figure 29 Summary of responses to Question 18

Comment: A wide range of the minimum amount of disk space requirements for production jobs were observed from the collected responses.

Question 19: To try and assess the requirements for the PRACE systems quantitatively, we would like you to score the following architecture features in terms of importance to your code. A total of 20 points should be distributed amongst the following requirements, with higher priority requirements receiving a higher number of points. If features are not important at all a score of zero can be used.

Response rate: 34/62

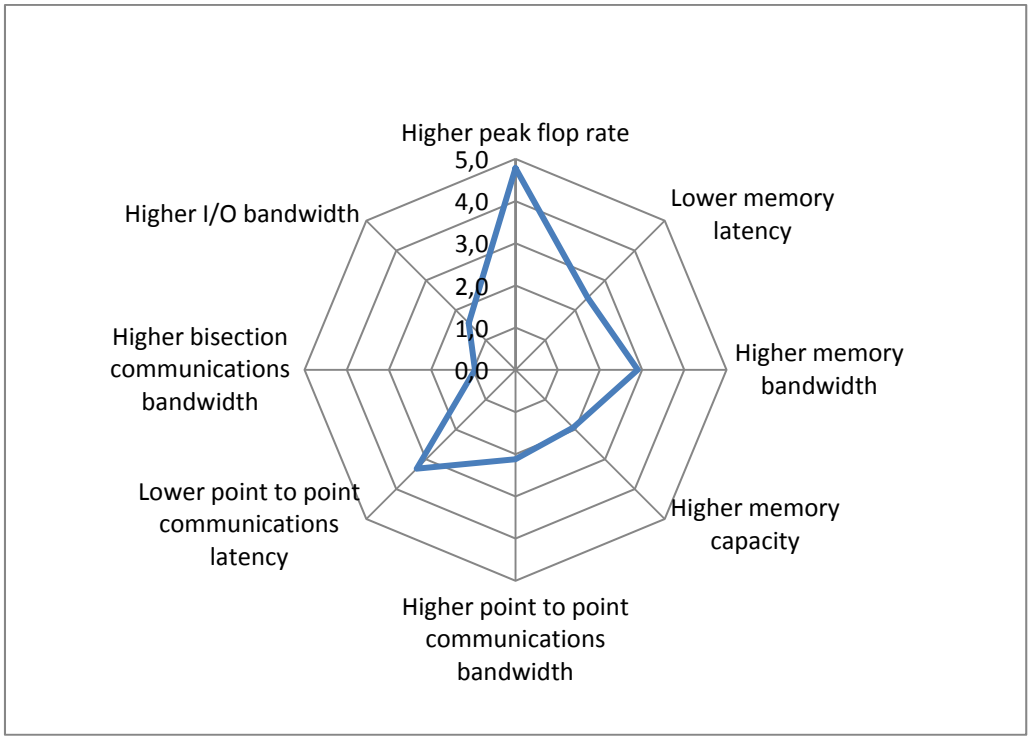

Figure 30 Summary of responses to Question 19

Comment: All listed features were considered important to some extent. “Higher peak flop rate” scored the highest, followed by “Lower point to point communications latency” and “Higher memory bandwidth”. “Higher bisection bandwidth” (the minimum bandwidth between two halves of the system, considering all possible ways to split the nodes into two equal partitions) scored lower than all the other features. The responses to this question are very consistent with responses to the same question in a previous survey of European HPC users [5].

Question 20: Are there any other architectural features that might affect the performance of your application?

Response rate: 9/62

Comment: Responses included higher collective communication performance and larger NUMA regions.

Question 21: Could your application benefit from accelerator devices, such as GPGPUs? (Single choice)

Response rate: 39/62

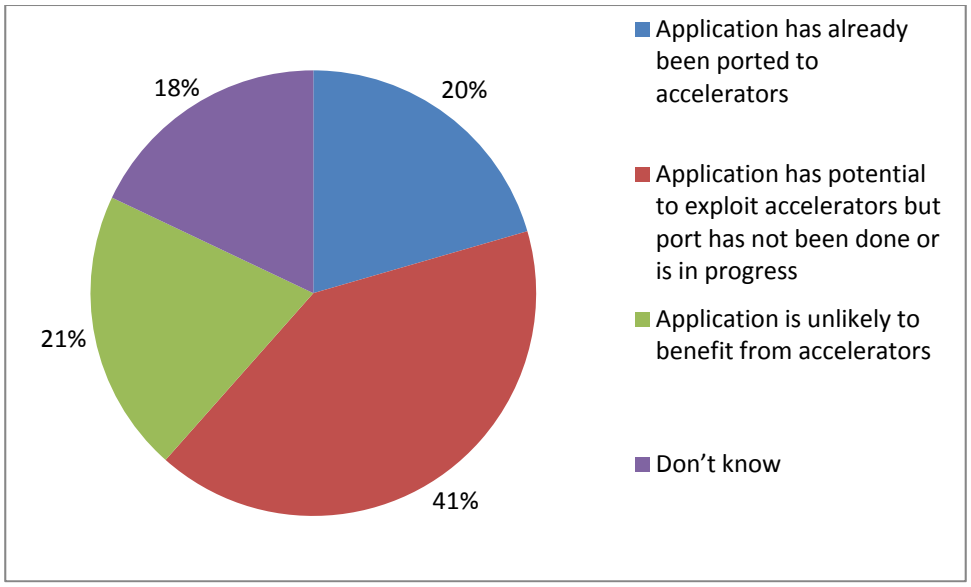

Figure 31 Summary of responses to Question 21

Comment: Around 61% of the responses indicated the applications have accelerator implementations or may potentially benefit from accelerators.

Question 22: What requirements do you have for network bandwidth between PRACE systems and other sites? (Multiple choice)

Response rate: 39/62

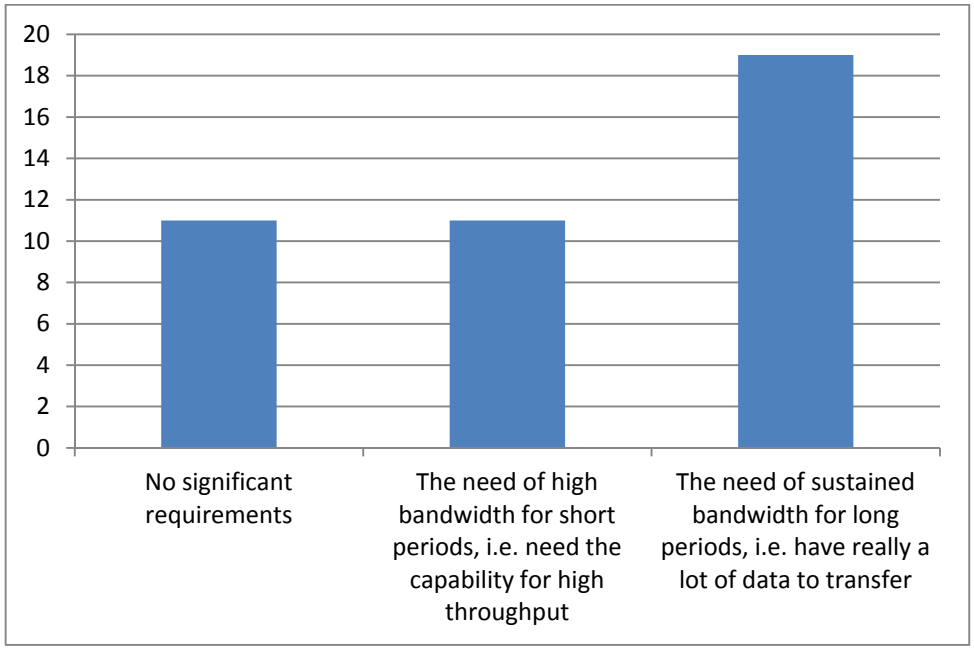

Figure 32 Summary of responses to Question 22

Comment: Both high throughput and large data transfer are important to a significant number of users.

Question 23: Do you know how to use the dedicated PRACE network, i.e. you know which systems are connected and which IP addresses to use? (Single choice)

Response rate: 39/62

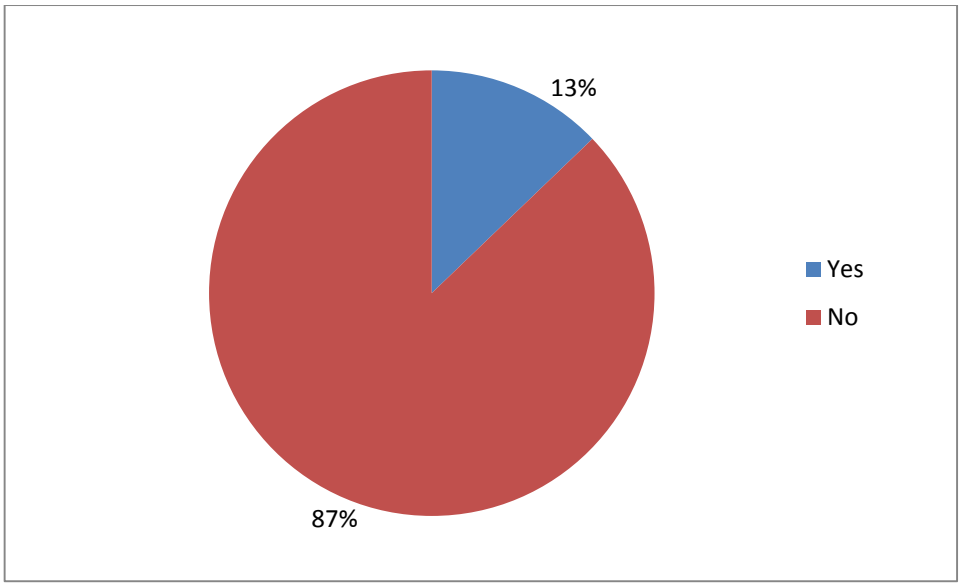

Figure 33 Summary of responses to Question 23

Comment: The responses suggest that awareness of this facility is low, but this may be because most of the survey responders only had access to one Tier-0 system.

Question 24: Which kind of additional network monitoring information are you interested in? (Multiple choice)

Response rate: 19/62

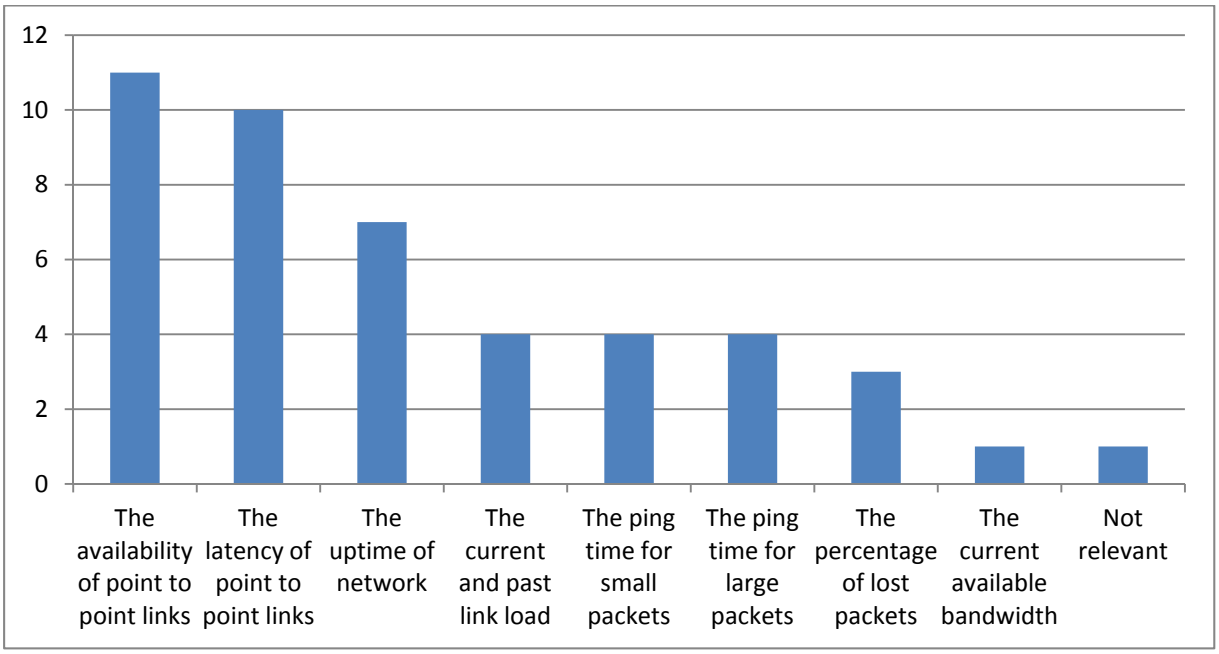

Figure 34 Summary of responses to Question 24

Comment: The responses to this question suggest a significant interest in having various types of network information available.

Question 25: Have you ever used the GridFTP service? (Single choice)

Response rate: 37/62

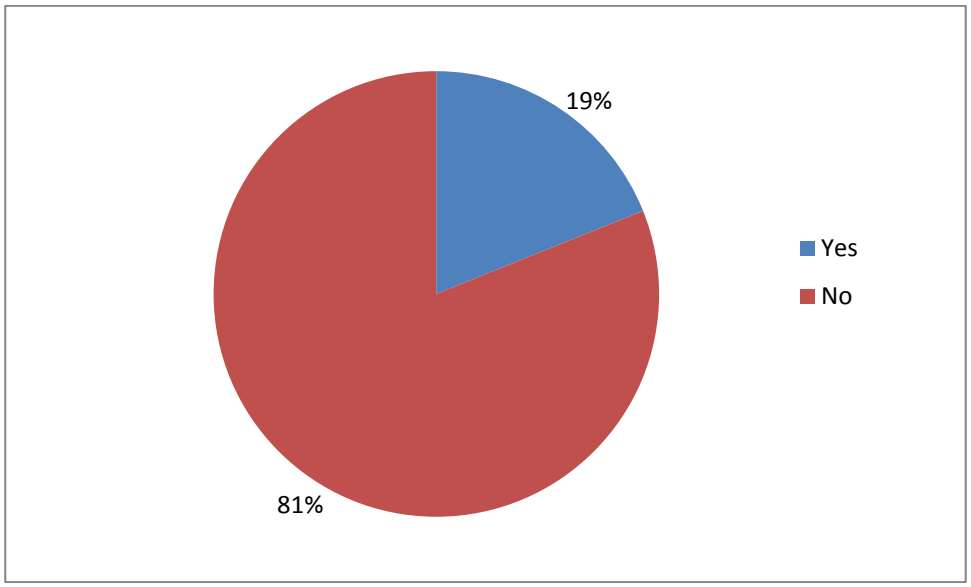

Figure 35 Summary of responses to Question 25

Question 26: Did you find GridFTP well suited to your requirements? (Single choice)

Response rate: 36/62

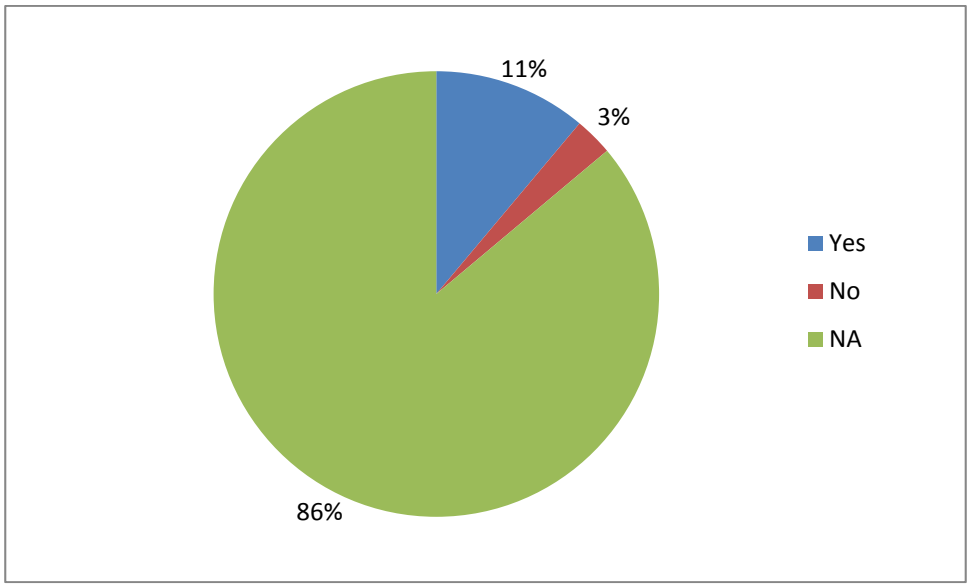

Figure 36 Summary of responses to Question 26

Question 27: Which job submission method do you use? (Single choice) &
Question 28: If you answered "Both" to the Question 27, please can you give the usage rate of the Local Batch System and UNICORE?

Response rate: 39/62 to Question 27; 0/62 to Question 28

Comment: All response to this question indicated that the job submission method used was Local Batch System (e.g. using batch commands like qsub, llsubmit, etc). There were no responses selecting "UNICORE" or "BOTH".

Question 29: If the following advanced features were available on PRACE Tier-0 systems, how often would you make use of them?

Response rate: 37/62

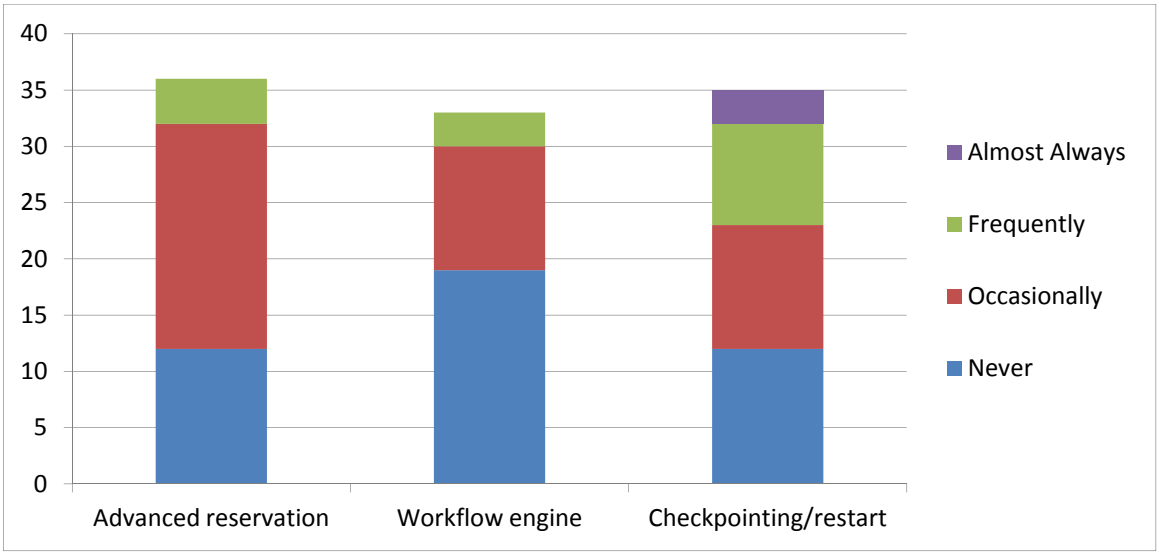

Figure 37 Summary of responses to Question 29

Comment: Besides the advanced features listed, one additional response mentioned the requirement for frequent usage of visualization facilities.

Question 30: Are you aware of the advantages of using Grid technology (Globus, UNICORE) for job submission or do you need Grid training? (Single choice)

Response rate: 37/62

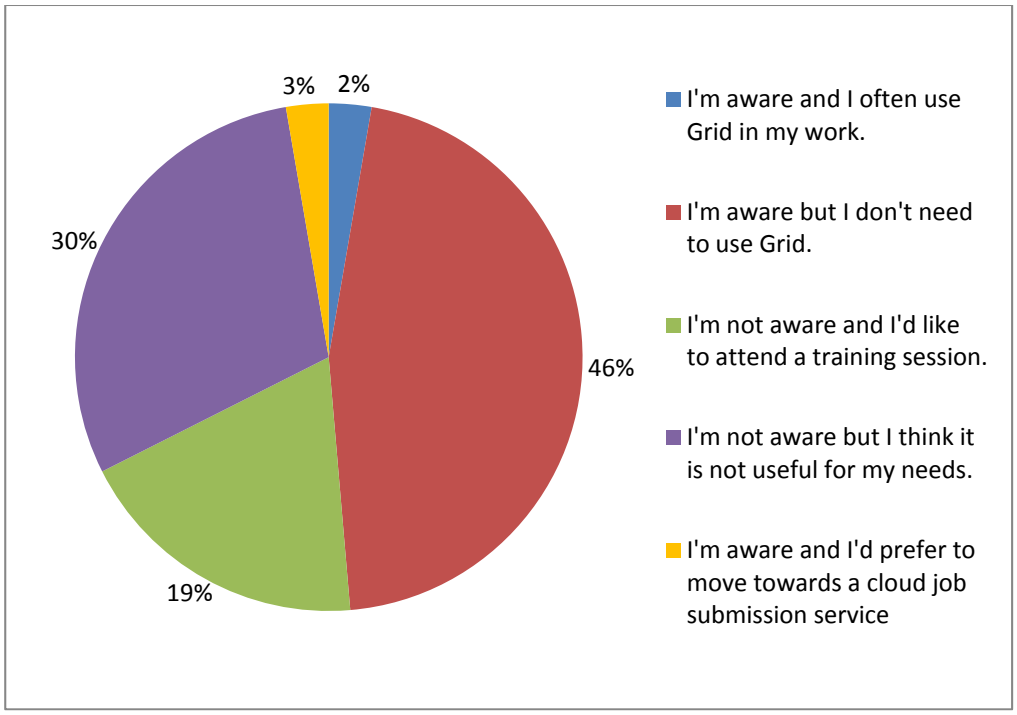

Figure 38 Summary of responses to Question 30

Comment: The responses suggest a low uptake of Grid technology.

Question 31: Which of the following Local Batch Systems/Schedulers are you familiar with? (Multiple choice)

Response rate: 39/62

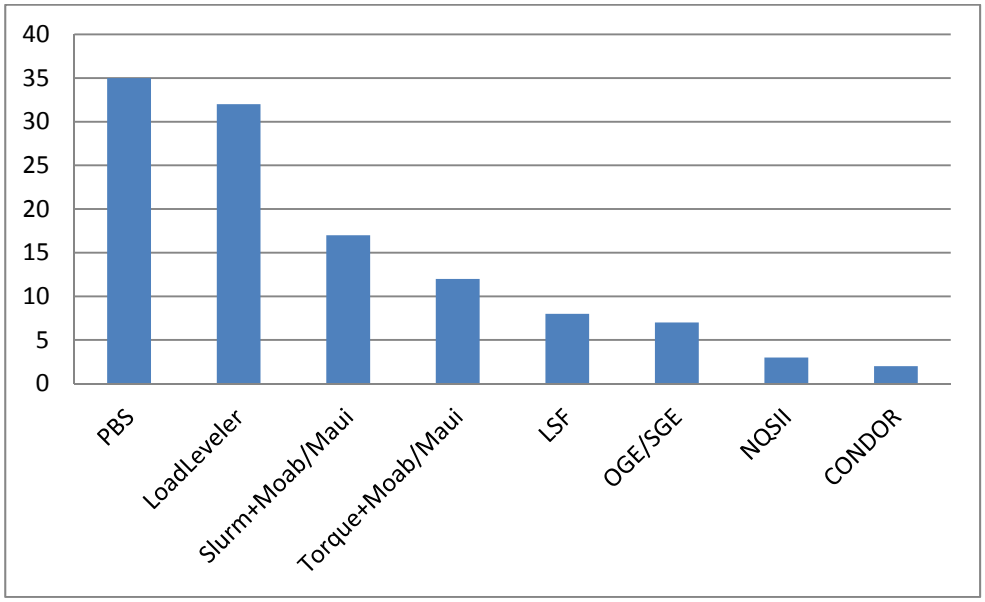

Figure 39 Summary of responses to Question 31

Question 32: Do you use customised/specific middleware to manage your jobs? (Single choice) &
Question 33: If you answered "Yes" to the previous Question 32, please specify which middleware you use for managing your jobs.

Response rate: 39/62 for Question 32, 4/62 for Question 33

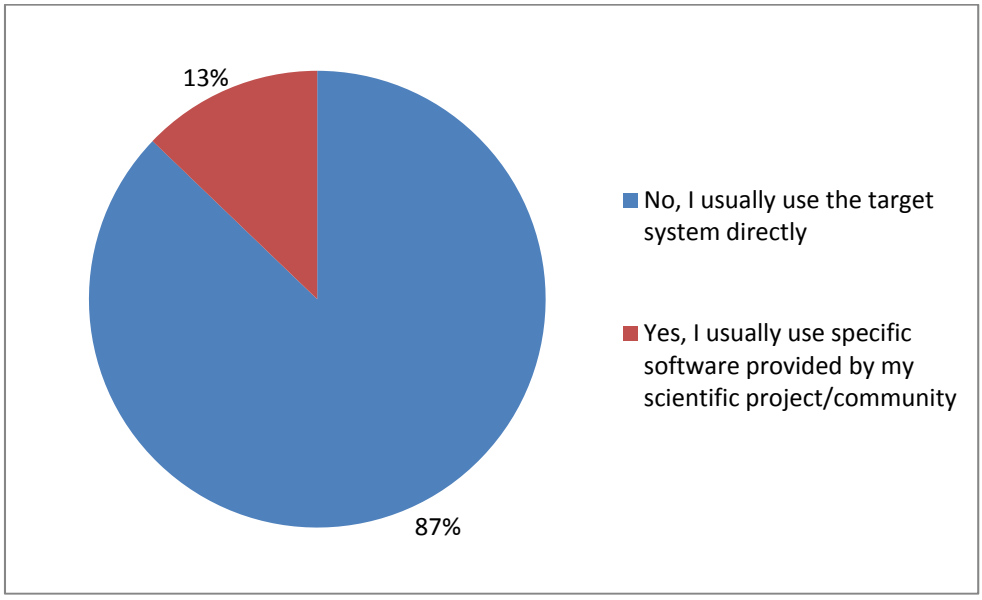

Figure 40 Summary of responses to Question 32

Comment: Only four responses were received for Question 33 and the middleware used were libIGCM, gLite and Python.

Question 34: Are you familiar with X.509 certificate-based authentication? (Single choice)

Response rate: 37/62

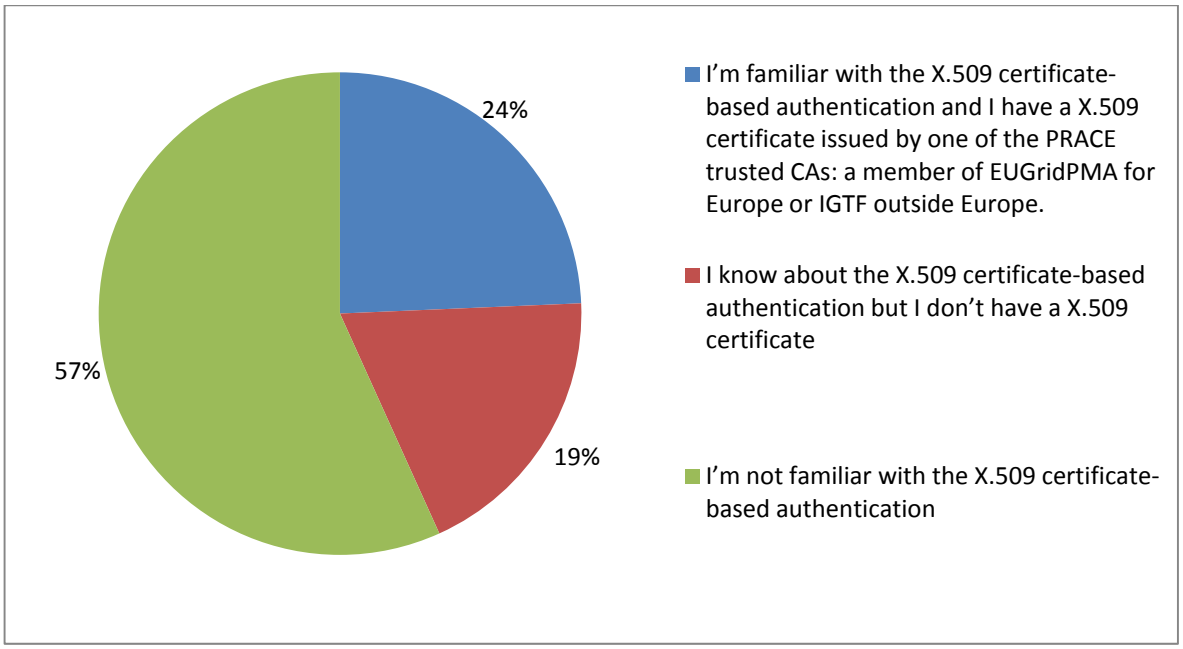

Figure 41 Summary of responses to Question 34

Question 35: Did you use your X.509 certificate during your work within PRACE access / project? (Single choice) &
Question 36: If you answered "No" to the Question 35, please let us know the reason why you didn't use the X.509 certificate for your PRACE access / project.

Response rate: 35/62 for Question 35, 16/62 for Question 36

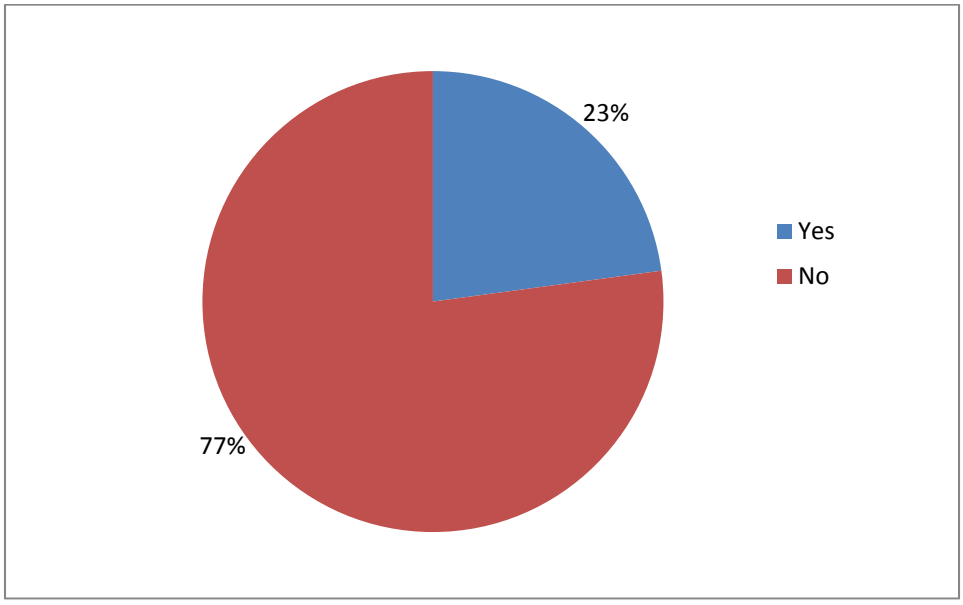

Figure 42 Summary of responses to Question 35

Comment: The reasons given in the responses to Question 36 are either because the users were not familiar with the X.509 certificate, or there was no need to use it.

Question 37: How often do / did you check the accounting information for your project? (Single choice)

Response rate: 39/62

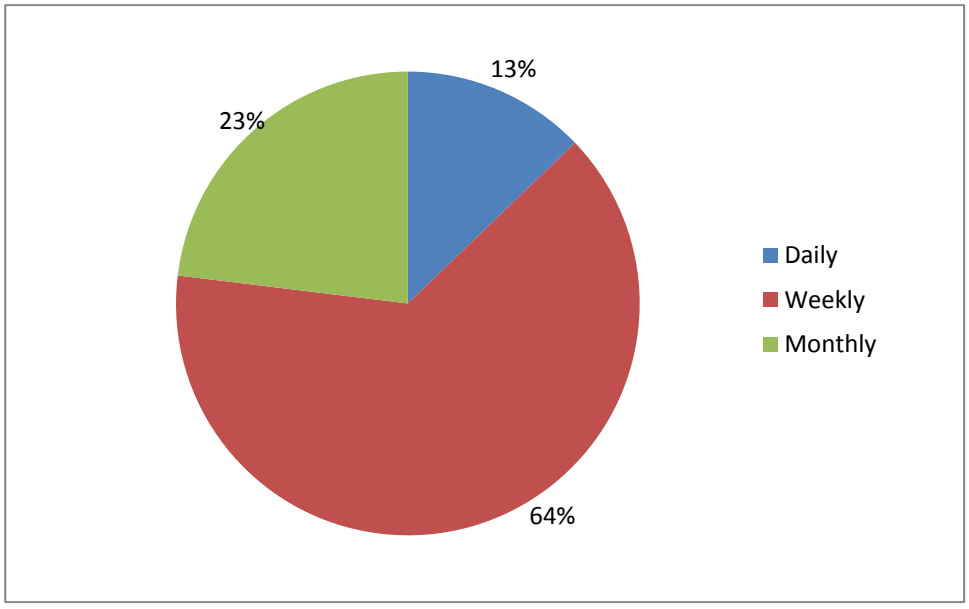

Figure 43 Summary of responses to Question 37

Question 38: Have you ever attend any PRACE training events? (Single choice)

Response rate: 40/62

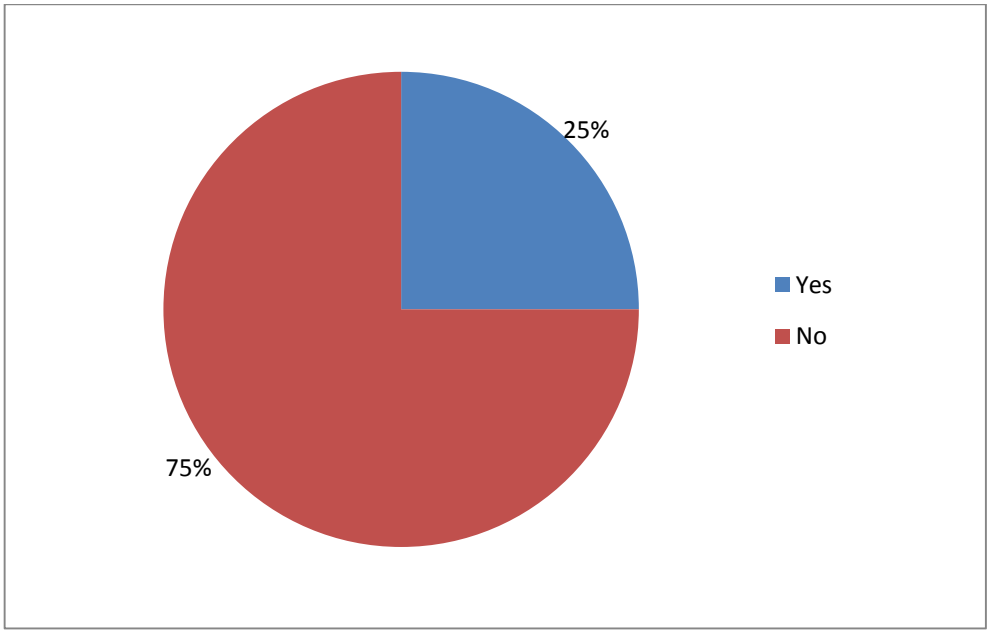

Figure 44 Summary of responses to Question 38

Question 39: Have you used any of the following ways to contact the PRACE site operator / host organisation in case of the problems with unexpected application behaviour? (Multiple choice)

Response rate: 40/62

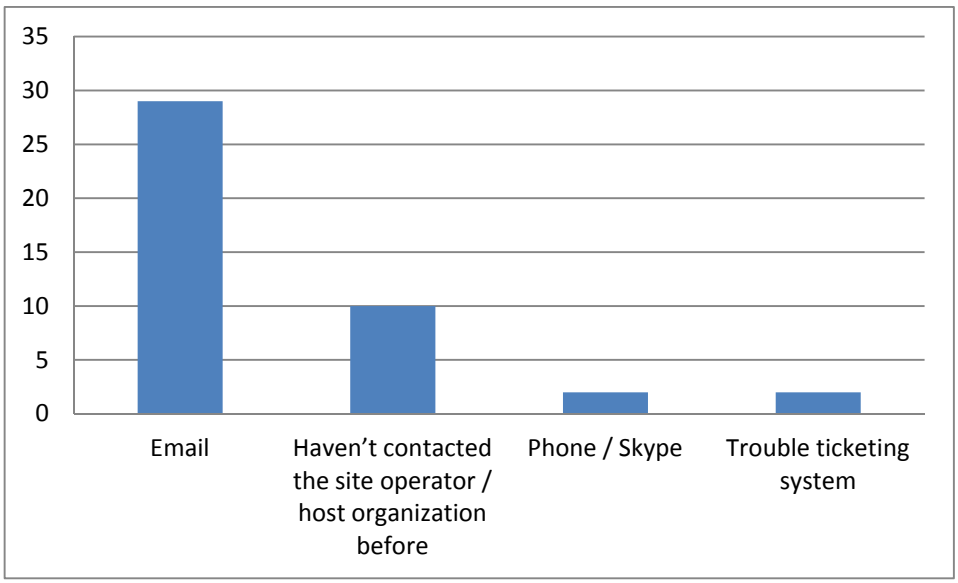

Figure 45 Summary of responses to Question 39

Question 40: During the lifetime of your project did you look up the following information?

Response rate: 36/62

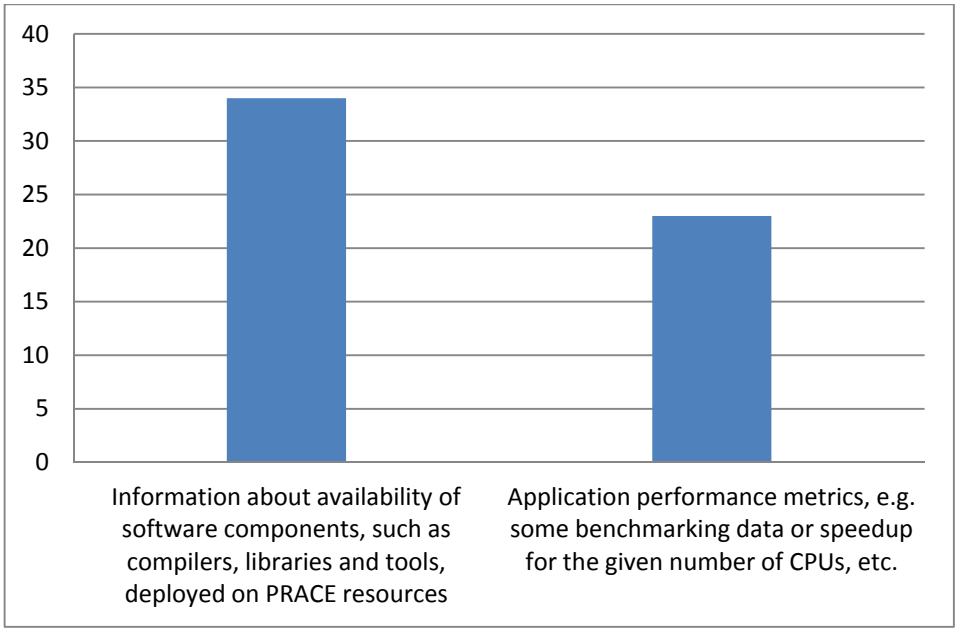

Figure 46 Summary of responses to Question 40

Comment: Both types of information are useful to users.

Question 41: What of the following would you like to access? (Multiple choice)

Response rate: 35/62

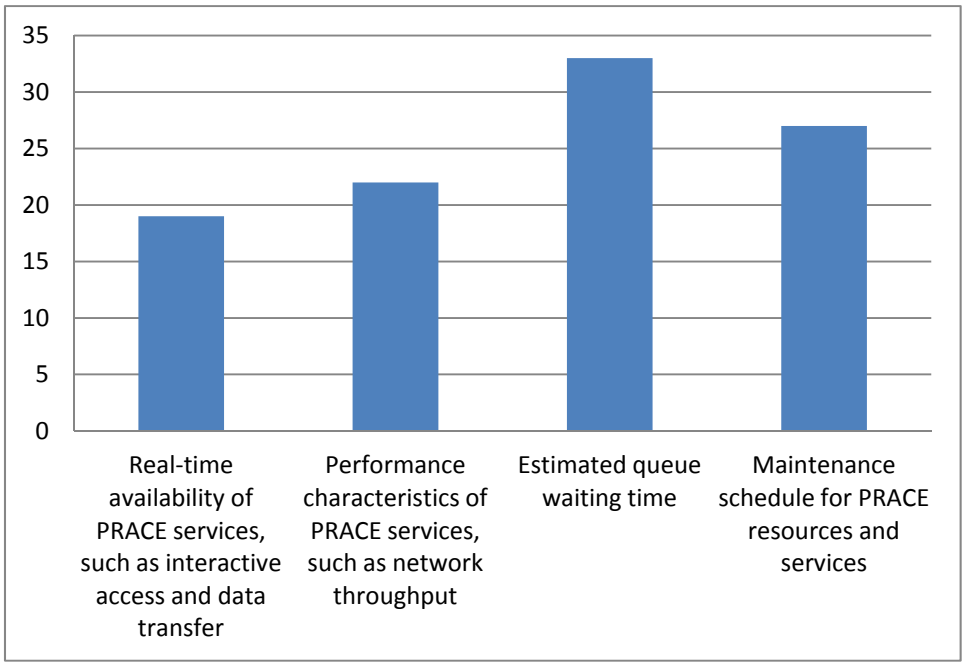

Figure 47 Summary of responses to Question 41

Comment: The responses suggest that all these types of information would be accessed by a significant number of users.

Question 42: Have you ever used/deployed interactive, 3D remote visualisation services / applications? (Single choice)

Response rate: 40/62

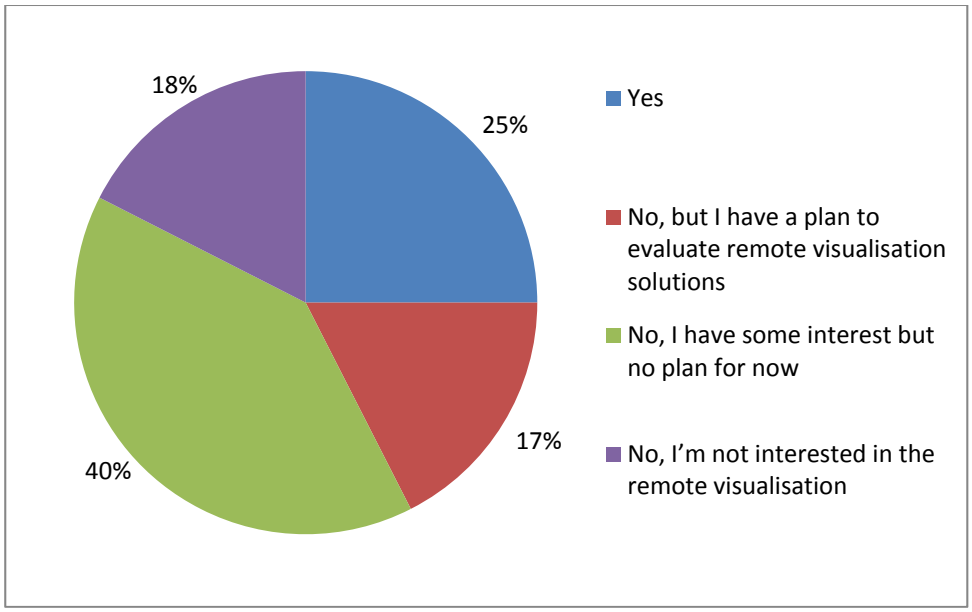

Figure 48 Summary of responses to Question 42

Question 43: If you have used / deployed 3D interactive applications using remote rendering through a VNC-like session, which solution did you use?

Response rate: 8/62

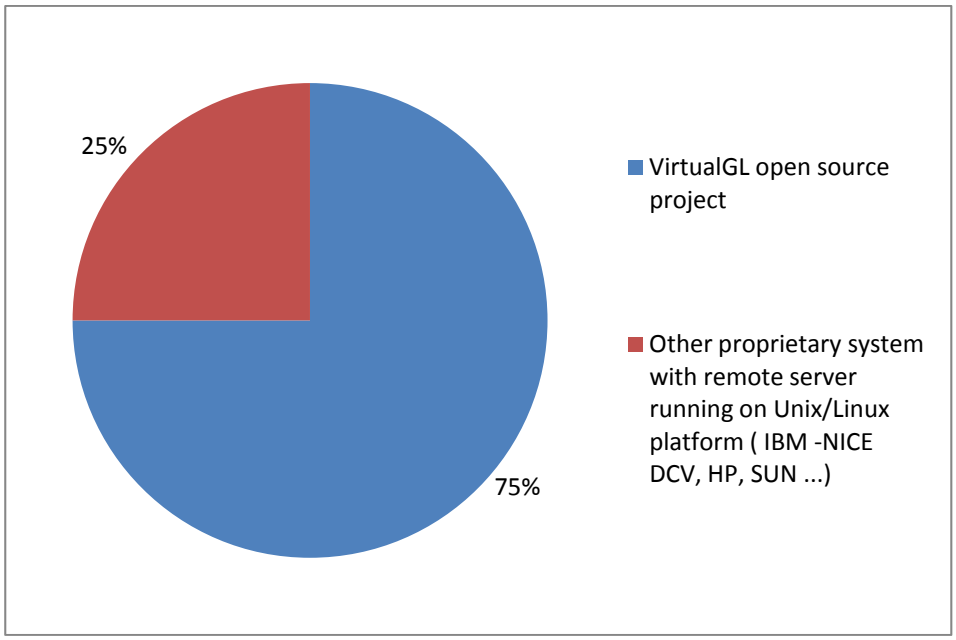

Figure 49 Summary of responses to Question 43

Question 44: If you have used / deployed applications with a client / server approach, exploiting server side (3D) visualisation, which solution did you use?

Response rate: 9/62

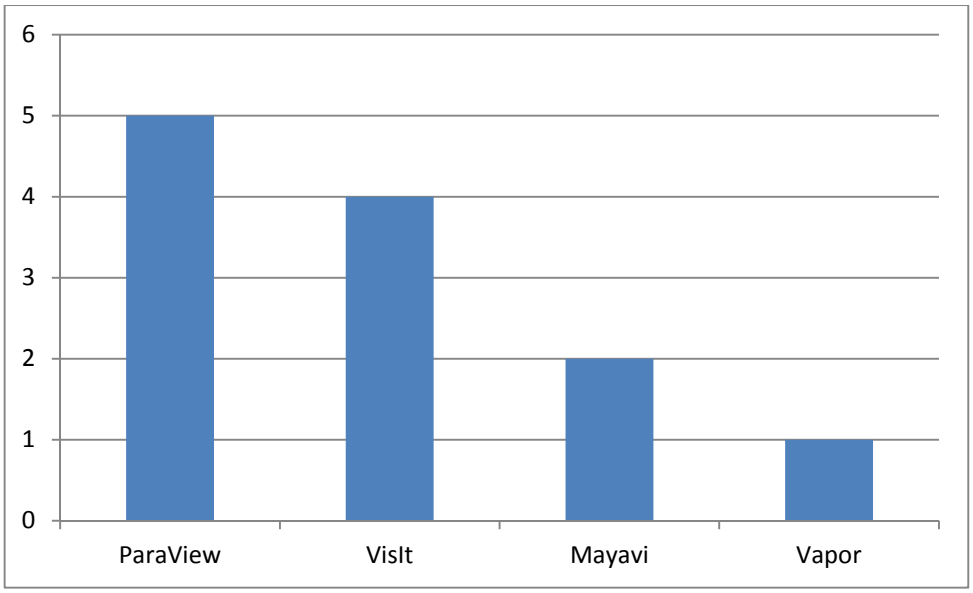

Figure 50 Summary of responses to Question 44

Question 45: Which visualisation remote infrastructure did you use / deploy?

Response rate: 15%

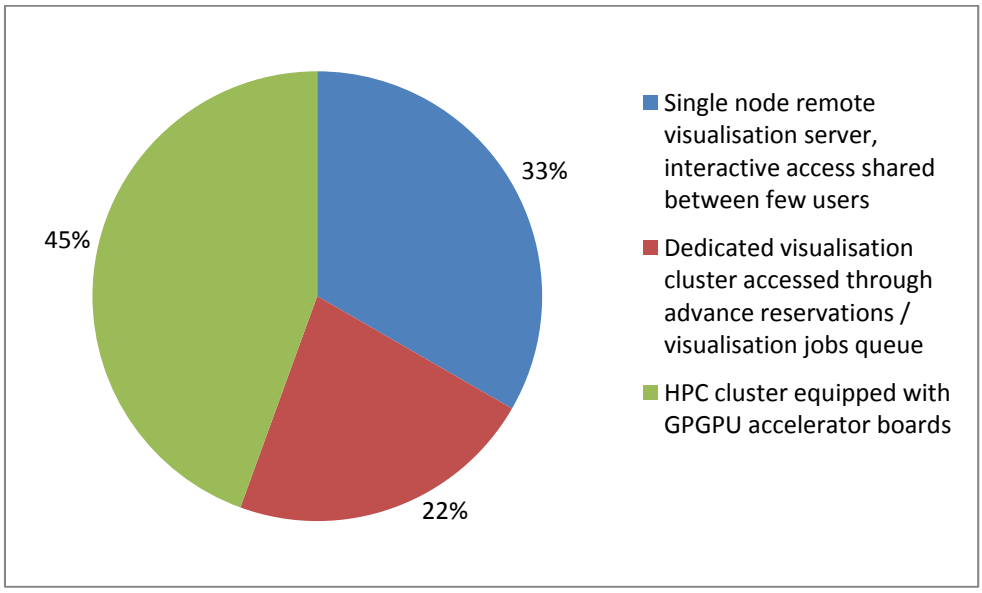

Figure 51 Summary of responses to Question 45

Question 46: What was the order of magnitude of the data size involved in one visualization session?

Response rate: 10/62

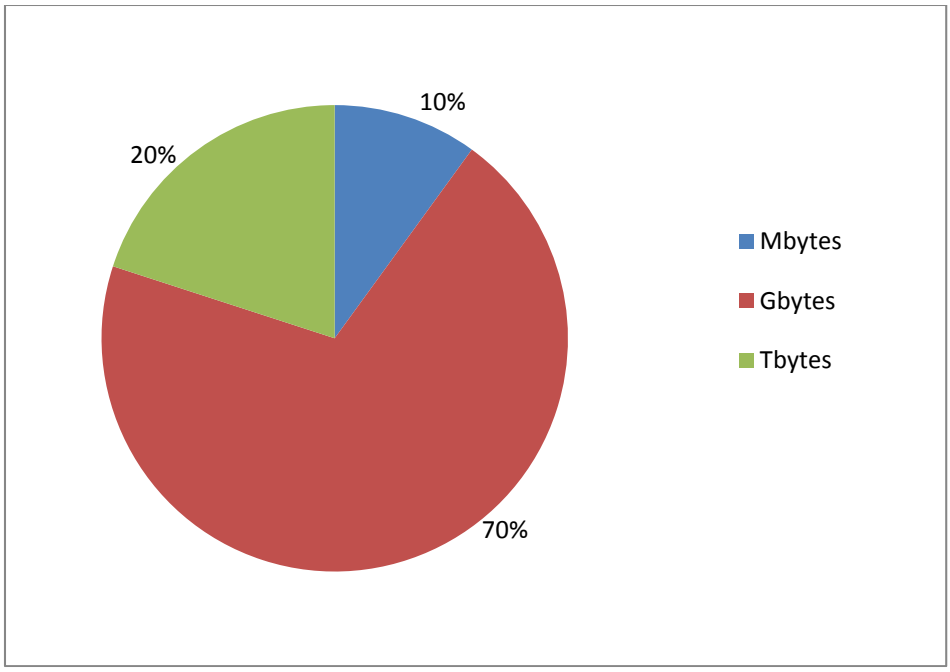

Figure 52 Summary of responses to Question 46

Question 47: What interconnect was used between your client platform and your remote visualization end?

Response rate: 9/62

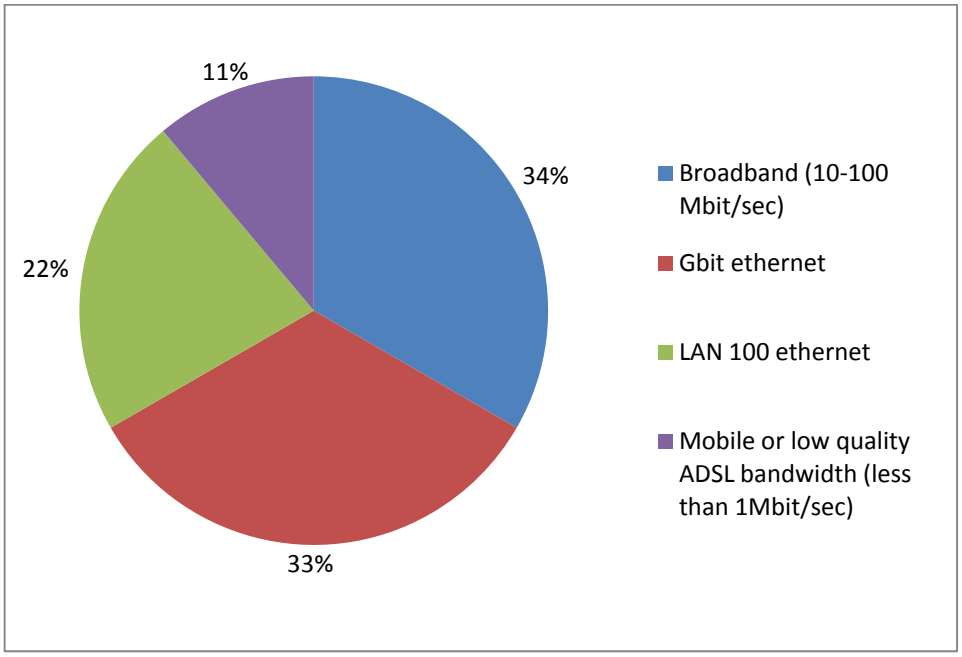

Figure 53 Summary of responses to Question 47

Question 48: The information gathered via this questionnaire will be used in a public deliverable. If there are any data protection policies or confidentiality issues, please state them here.

Response rate: 0/62

Comment: No confidentiality concerns were identified.

3.3 Summary

We have conducted a survey of users participating in PRACE Preparatory and Regular Access projects. The principal findings from this survey are:

- Nearly 50% of the application codes use combined MPI+OpenMP as the parallelisation method. (The remainder use MPI only: no use was reported of other parallelization models such as PGAS.)
- There is a roughly equal split between Fortran and C/C++ as the main language used by applications.
- Higher peak flop rate is seen as the architectural feature that would most benefit applications.
- Over 60% of the applications were reported as having an existing accelerator port, or would be likely to benefit from such.
- In terms of network requirements between PRACE systems and other sites, both capability and capacity are seen as important.
- Only 13% of users know how to use the dedicated PRACE network.
- There is significant interest in having various types of network information available to users.
- Use of, and interest in, Grid technologies is low.
- A significant minority of users is interested in, or makes use of, remote visualization facilities.

4 Conclusions and future work

We have carried out of PRACE usage of Tier-0 systems, and of users participating in PRACE Preparatory and Regular Access projects. These surveys have received a good response rate, and have generated a significant amount of data. This deliverable provides a top-level summary of the survey results. The data from these surveys are available to other tasks/work packages in PRACE for further analysis, as required.

From the survey of Tier-0 systems, we have found that good use is being made of these systems, in that CPU hours committed by hosting partners and allocated to PRACE Access projects are being fully consumed, and that almost all projects are running jobs on high numbers of CPU cores. Disk usage varies widely between projects, with a small number of projects consuming the majority of disk space.

From the survey of PRACE Access users, we have found that the roles of project manager, code developer and scientific end-user are all well represented amongst the users. With respect to the applications being used, there has been a significant increase in the number of applications which use hybrid MPI+OpenMP as the parallel programming model, compared to earlier surveys carried out by the PRACE project. Fortran and C/C++ are still the dominant programming languages, and are roughly equally popular among applications. There has been essentially no uptake of PGAS languages.

In agreement with previous surveys, users still view increased flop rate as the architectural enhancement that would most benefit the performance of their applications. Over half the applications mentioned were reported as having existing or in-progress accelerator ports, indicating that there may be significant demand for this style of architecture.

Awareness of the dedicated PRACE network is low, but this is unsurprising since this survey covered only two Tier-0 systems, and few users have accounts on both. Capacity and capability of networking are viewed as equally important, and users show significant interest in the availability of network information. Both the use of, and interest in Grid technologies is low, suggesting that the demand for this type of software is limited.

A significant minority of users reported that they used, or were interested in, remote visualization facilities.

It is intended that this type of survey of user experiences and requirements be repeated in the future, which will keep the information up to date and to allow trends and changes to be observed. Task 6.3 in PRACE 3IP “Technical evolution of the PRACE operational services” will be responsible for carrying out this work.